Release Notes: Junos® OS Release 20.1R3 for the ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR

12 May 2022

Contents

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, EX Series, JRR Series, Junos fusion for enterprise, Junos fusion for provider edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

These release notes accompany Junos OS Release 20.1R3 for the ACX Series, EX Series, JRR Series, Junos fusion for enterprise, Junos fusion for provider edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX and vRR. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

- **In Focus guide**—We have a document called In Focus that provides details on the most important features for the release in one place. We hope this document will quickly get you to the latest information about Junos OS features. Let us know if you find this information useful by sending an e-mail to techpubs-comments@juniper.net.
- **Important Information:**
  - Upgrading Using ISSU on page 366
  - Licensing on page 366
  - Compliance Advisor on page 366
  - Finding More Information on page 367
  - Documentation Feedback on page 367
  - Requesting Technical Support on page 369

Junos OS Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 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 Release 20.1R3

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 20.1R3.

What's New in Release 20.1R2

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 20.1R2.

What's New in Release 20.1R1

Interfaces and Chassis

- **Support for new show | display set CLI commands (ACX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the following new **show** commands have been introduced:
  - **show | display set explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the top level of the hierarchy.
  - **show | display set relative explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the current hierarchy level.

  [See show | display set and show | display set relative.]

Junos OS XML API and Scripting

- **The jcs:load-configuration template supports loading the rescue configuration (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **jcs:load-configuration** template supports the **rescue** parameter to load and commit the rescue configuration on a device. SLAX and XSLT scripts can call the **jcs:load-configuration** template with the **rescue** parameter set to **"rescue"** to replace the active configuration with the rescue configuration.

  [See Changing the Configuration Using SLAX and XSLT Scripts and jcs:load-configuration Template.]

Junos Telemetry Interface

- **gRPC Dial-Out support on JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Junos OS Release 20.1R1 provides remote procedure call (gRPC) dial-out support for telemetry. In this method, the target device (server) initiates a gRPC session with the collector (client) and, when the session is established, streams the telemetry data that is specified by the sensor-group subscription to the collector. This is in contrast to the gRPC network management interface (gNMI) dial-in method, in which the collector initiates a connection to the target device.

  gRPC dial-out provides several benefits as compared to gRPC dial-in, including simplifying access to the target advice and reducing the exposure of target devices to threats outside of their topology.

  To enable export of statistics, include the **export-profile** and **sensor** statements at the [**edit services analytics**] hierarchy level. The export profile must include the reporting rate, the transport service (for example, gRPC), and the format (for example, gbp-gnmi). The sensor configuration should include the name of the collector (the server's name), the name of the export profile, and the resource path. An example of a resource path is **/interfaces/interface[name='fxp0'**.

  [See Using gRPC Dial-Out for Secure Telemetry Collection.]

- **gRPC version v1.18.0 with JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Junos OS Release 20.1R1 includes support for remote procedure call (gRPC) services version v1.18.0 with Junos telemetry

interface (JTI). This version includes important enhancements for gRPC. In earlier releases, JTI is supported with gRPC version v1.3.0.

Use gRPC in combination with JTI to stream statistics at configurable intervals from a device to an outside collector.

[See gRPC Services for Junos Telemetry Interface.]

MPLS

- **CoS-based forwarding and policy-based routing to steer selective traffic over an SR-TE path (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 20.1R1, you can use CoS-based forwarding (CBF) and policy-based routing (PBR, also known as filter-based forwarding or FBF) to steer service traffic using a particular segment routing-traffic-engineered (SR-TE) path. This feature is supported only on uncolored segment routing LSPs that have the next hop configured as a first hop label or an IP address.

  With CBF and PBR, you can :

  - Choose an SR-TE path on the basis of service.
  - Choose the supporting services to resolve over the selected SR-TE path.

[See Example: Configuring CoS-Based Forwarding and Policy-Based Routing For SR-TE LSPs.]

Routing Protocols

- **Support for topology-independent loop-free alternate (TI-LFA) in IS-IS for IPv6-only networks (ACX Series, MX Series, and PTX Series)**— Starting with Junos OS Release 20.1R1, you can configure TI-LFA with segment routing in an IPv6-only network for the IS-IS protocol. TI-LFA provides MPLS fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. TI-LFA provides protection against link failure, and node failure.

  You can enable TI-LFA for IS-IS by configuring the **use-post-convergence-lfa** statement at the **[edit protocols isis backup-spf-options]** hierarchy level. You can enable the creation of post-convergence backup paths for a given IPv6 interface by configuring the **post-convergence-lfa** statement at the **[edit protocols isis interface *interface-name* level *level*]** hierarchy level. The **post-convergence-lfa** statement enables link-protection mode.

  You can enable **node-protection** mode for a given interface at the **[edit protocols isis interface *interface-name* level *level* post-convergence-lfa]** hierarchy level. However, you cannot configure fate-sharing protection for IPv6-only networks.

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

System Management

- **Restrict option under NTP configuration is now visible (ACX Series, QFX Series, MX Series, PTX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **noquery** command under the **restrict** hierarchy is now available and can be configured with a mask address. The **noquery** command is used to restrict ntpq and ntpdc queries coming from hosts and subnets.

  [See Configuring NTP Access Restrictions for a Specific Address.]

SEE ALSO

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 20.1R3

General Routing

- **cRPD supports the Junos Telemetry Interface (JTI) over TLS similar to Junos OS (cRPD)**—cRPD supports local (server-side) certificate validation for gRPC and JTI similar to Junos OS. cRPD doesn't support bidirectional authentication for gRPC and JTI.

  [See Configuring gRPC for the Junos Telemetry Interface and Importing SSL Certificates for Junos XML Protocol Support.]

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.]

- **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).]

MPLS

- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the cspf-backoff-time statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

  [See cspf-backoff-time.]

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.]

Routing Protocols

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the lsdist.0 and lsdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, we added multiple secondary loopback addresses in the traffic engineering database to the lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router ID.

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.]

What's Changed in Release 20.1R2

Class of Service (CoS)

- We've corrected the output of the **show class-of-service interface | display xml** command. The output is of the following sort: <container> <leaf-1> data </leaf-1><leaf-2>data </leaf-2> <leaf-3> data</leaf-3> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container> will now appear correctly as <container> <leaf-1> data </leaf-1><leaf-2>data </leaf-2> <leaf-3> data</leaf-3></container> <container> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container>.

General Routing

- **Support for gigether-options statement (ACX5048 and ACX5096)**— Junos OS supports the **gigether-options** statement at the **edit interfaces interface-name** hierarchy on the ACX5048 and ACX5096 routers. Previously, support for the **gigether-statement** was deprecated.

  [See gigether-options .]

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**— You can set the verbosity of the trace log to only show error messages using the error option at the **edit system services extension-service traceoptions level** hierarchy.

  [**See** traceoptions (Services).]

What's Changed in Release 20.1R1

There are no changes in behavior and syntax for ACX Series in Junos OS Release 20.1R1.

SEE ALSO

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 logical interface-level statistics do not get the expected rate PPS values with the traffic flowing. PR1217154
- Filter for sensors (Optical and OTN) does not work. PR1371996
- On the ACX5448-D interface, when configured over Layer 2 VPN with control-channel set, BFD gets to the **Down** state. PR1432854
- The time consumed on 1-Gigabit Ethernet performance is not the same compared to 10-Gigabit Ethernet. Compensation is done to bring the mean value under class A but the peak-to-peak variations are high and might go beyond 100 ns. It has a latency variation with peak-to-peak variations of around 125 ns—250 ns without any traffic. (For example, 5-10 percent of the mean latency introduced by each phy, which is of around 2.5 microseconds). PR1437175
- With an asymmetric network connection, a 10-Gbps MACsec port connected to a 10-Gbps channelized port, high and asymmetric T1 and T4 time errors are observed. This situation introduces a high two-way time error and also different CF updates in the forward and reverse paths. PR1440140
- With the MACsec feature enabled and introduction of traffic, the peak-to-peak value varies with the percentage of traffic introduced. Finding the maximum and mean values of the time errors with different traffic rates (for example, two-router scenarios) can have the maximum value as high as 1054 ns with 95 percent traffic, 640 ns for 90 percent traffic, and 137 ns with no traffic. PR1441388
- IGMPv3 rate of join supported in the ACX5448 routers or an ACX5448-D interface is around 900 joins per second. PR1448146
- On the ACX5448 routers, the OSPF state is not as expected.PR1543667

SEE ALSO

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

- In ACX routers acting as PE routers for L2VPN/L2circuit/VPLS/L3VPN services, traffic forwarding over the MPLS paths used by the VPLS/L2VPN/L2circuit/L3VPN services can be stopped. The traffic forwarding issue happens when the LSP flaps between primary and backup paths in a particular sequence. PR1204714
- Loopback status is not shown for OT interfaces on CLI (available from vty only). PR1358017
- The SD (Signal Degrade) threshold is normally lower than the SF threshold (so that as errors increase, SD condition is encountered first). For the ACX6360 optical links there is no guard code to prevent the user from setting the SD threshold above the SF threshold which would cause increasing errors to trigger the SF alarm before the SD alarm. This will not cause any issues on systems with correctly provisioned SD or SF thresholds. PR1376869

- On the ACX6360-OR router, enhancement is needed for the FRR BER threshold SNMP support. PR1383303
- On the ACX6360 router, Tx power cannot be configured using + sign.PR1383980
- The **ccc** logs are not compressed after rotation. PR1398511
- Link fault signaling (LFS) feature is not supported on ACX5448 10/40/100GbE interfaces. PR1401718
- A jnxIfOtnOperState trap notification is sent for all OT interfaces. PR1406758
- When a timing configuration and the corresponding interface configuration are flapped for multiple times in iteration, PTP is stuck in "INITIALIZE" state where the ARP for the neighbor is not resolved. In issue state, BCM hardware block get into inconsistency state, where the lookup is failing. PR1410746
- On ACX5048 and ACX5096 platforms, traffic loss and SNMP slow response issues could be seen where an optic transceiver is removed and inserted back to the same interface. Manually restarting Packet Forwarding Engine might also trigger this defect. PR1418696
- On an ACX5000 platform, high CPU usage by the fxpc process might be seen under a rare condition if parity errors are detected in devices. This issue has no direct service/traffic impact. However, because CPU utilization is high during this issue, there are some side effects. For example, the issue could impact time-sensitive features such as BFD. PR1419761
- The em2 interface configuration causes FPC to crash during initialization and FPC does not come online. After deleting the em2 configuration and restarting the router, FPC comes online. PR1429212
- Protocols get forwarded when using a non-existing SSM map source address in IGMPv3 instead of pruning. PR1435648
- Memory leaks are expected in this release. PR1438358
- On an ACX5448 box, link flaps or CoS configuration changes (specific to temporal value changes) might result in traffic drop on all interfaces and recorded as RED drops. PR1443466
- Recovery of Junos OS volume from OAM becomes nonresponsive indefinitely. PR1446512
- Drop profile maximum threshold might not reach its full limit when the packet size is other than 1000 bytes. PR1448418
- The CFM REMOTE MEP does not come up after configuration or if the MEP remains in the Start state. PR1460555
- On an ACX710 device, MPLS packet load balancing is done without hashing enabled.PR1475363
- FPC might continuously crash after deactivating or activating loopback filter or reboot the system after configuring the loopback filter. PR1477740
- On the ACX5048 router, the fxpc process generates core files after ISSU is terminated. PR1489765
- On the ACX5000 router, the IEEE 802.1p priority and DEI values in the locally generated VLAN-based IP packets might be changed when sourced from the IRB interface. PR1490966

- In PTP environment some vendor devices acting as slave expecting announce messages at an interval of -3 (8pps) from upstream master device. As of today announce message are configurable in range of 0 to 3. To support the above requirement engineering provided a hidden cli configuration **set protocol ptp master announce-interval -3**. In the networks/design where we have this requirement we can configure the hidden cli otherwise regular cli which is in the range (0 to 3) can be configurable. Both the cli configurations are mutually exclusive, commit error is expected if both are configured. This new change is applicable to ACX platforms only excluding ACX5000. PR1507782
- On the ACX5448 routers, the transit DHCP packets drop is observed.PR1517420
- On the ACX5448 routers, the OSPF state is not as expected.PR1543667
- Even though enhanced-ip is active, the following alarm is observed during ISSU: **RE0 network-service mode mismatch between configuration and kernel setting**. PR1546002
- As per the current code, ACX would not delete a mac address from the mac table there is- (a) traffic destined to the mac address or (b) traffic sourced from the mac address or (c) both Fix of this PR will allow ACX to only look at (b) traffic sourced from mac address before deleting the mac address entry from mac table. So, if there is no traffic sourced from the mac for an interval of mac aging timer, the mac would be deleted from the mac table at the end of mac aging timer with out taking into account the traffic destined to the mac address. PR1565642
- The Precision Time Protocol (PTP) clock might fail to lock and stuck in acquiring state at clock servo. PR1570310
- If the BFD session is configured with minimum-interval as 4ms, due to HW issue, minimum interval is not configured as expected. So, the default minimum interval will get applied for the session. PR1585382
- On deleting remote-mep and auto-discovery in the peer, RDI sent by peer device is not processed in DUT. Hence, action configured for RDI event in the action-profile is not happening as expected. PR1592571

Network Management and Monitoring

- On all Junos platforms, the SNMP polling might not be working if the ISIS protocol is disabled under the same VRF (Virtual Routing and Forwarding) through which SNMP requests are sent. PR1527251

Platform and Infrastructure

- A buffer overflow vulnerability in the TCP/IP stack of Juniper Networks Junos OS allows an attacker to send specific sequences of packets to the device thereby causing a Denial of Service (DoS). Please refer to https://kb.juniper.net/JSA11200 for more information. PR1557881

Virtual Chassis

- The ACX5000 reports false parity error messages such as **soc_mem_array_sbusdma_read**. The ACX5000 SDK can raise false alarms for parity error messages such as **soc_mem_array_sbusdma_read**. This is a false positive error message. PR1276970

SEE ALSO

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: 20.1R3

General Routing

- The IPv6 BFD sessions flap when configured below 100 ms flaps. PR1456237
- On ACX5448, the MAC learning and aging might not work if there are excessive MAC movements or continuous interface flaps. PR1480235
- On the ACX710 routers, the following error message is observed: **PFE_ERROR_FAIL_OPERATION: Failed to install in h/w, LOG: Err] dnx_nh_unilist_install: BCM L3 Egress create object failed for:Unilist nh 2097369 (0:Ok) nh 0.** PR1495563
- The hardware FRR for EVPN-VPWS, EVPN-FXC, and Layer 3 VPN with a composite next hop are not supported. PR1499483
- The ACX1100, ACX2100, ACX2200, ACX2000, and ACX4000 routers might stop forwarding transit and control traffic. PR1508534
- On the ACX500-I router, the **show services session count** command does not work as expected. PR1520305
- The interface does not come up with the autonegotiation setting between ACX1100 routers and QFX Series switches, MX Series routers, and ACX Series routers as the other end.PR1523418
- Packet drops might be seen with all commit events with 1G speed configured interface PR1524614
- The aggregated Ethernet interface might not come up with LFM configured after reboot. PR1526283
- With the ACX5448 router with 1000 CFM, the CCM state does not go in to the Ok state after loading the configuration or restarting the Packet Forwarding Engine. PR1526626
- The l2cpd memory leak might be observed with aggregated Ethernet interface flap. PR1527853
- Packets drops might be seen after configuring the PTP transparent clock. PR1530862
- The BUM (Broadcast, Unknown Unicast, and Multicast) traffic might drop in the VPLS instance under certain conditions. PR1531733
- VPLS traffic may be discarded on the ACX5448 and ACX710 products. PR1535250
- Snmp mib walk for jnxSubscriber OIDs returns a general error. PR1535754
- On the ACX5448 router, unexpected behavior of the **show chassis network-services** command is observed. PR1538869
- SFP-T interface might not come up if a straight cable is used in ACX5448. PR1547394
- The ACX5448 router as transit for the BGP-labeled unicast drops traffic. PR1547713

- On the ACX5048 router, the fxpc process generates the core file on the analyzer configuration. PR1559690
- The DF(Designated Forwarder) might not forward traffic. PR1567752
- ACX resets tunable optics to default wavelength after upgrade or reboot. PR1570192
- Packets might get tagged with the default VLAN-ID and dropped at the peer under Layer 2 circuits local switching scenario. PR1574623
- ACX as a LSR router, fails to process RSVP Path Message. PR1576585
- There might be a traffic drop between customer edge and provider edge devices in case of ARP resolution failure. PR1580782

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 PR1556103
- FPC crash might be observed after the **show class-of-service** command. PR1568661

Infrastructure

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

Layer 2 Features

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

Routing Protocols

- The **mld snooping membership** command is not accepting group with vlan options together (same works with IGMP snooping membership command). PR1516650
- The rpd memory leak might be seen in the BGP scenario. PR1547273

VPNs

- The l2circuit local-switching end interface might get stuck in XX (Unknown) state upon vlan-id-list configuration change. PR1528809

Resolved Issues: 20.1R2

Resolved Issues: 20.1R2

General Routing

- Policer discarded count is displayed incorrectly in the enqueue count of the interface queue, but the traffic behavior is as expected. PR1414887
- On the ACX2000 router, which is a PTP hybrid, drift messages are observed. PR1426910

- The **gigether-options** command is enabled again under the interface hierarchy. PR1430009
- The statistics are accessed through Broadcom API, which is the same for both tagged and untagged packets. This cannot be changed in accordance with the MX Series routers because the statistics are directly accessed from Broadcom without any statistics changes specific to tagging from the ACX5448 router side. This impacts other statistics if the change is made. PR1430108
- While powering off or powering on the device repeatedly, the SMBUS transactions timeout occurs. PR1463745
- Unable to get shared buffer count as expected. PR1468618
- The router might become nonresponsive and bring the traffic down when the disk space becomes full. PR1470217
- On the ACX5048 router, the egress queue statistics do not work for the aggregated Ethernet interfaces. PR1472467
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. PR1473610
- The links might not come up when the 100-Gigabit Ethernet interface is channelized into the four 25-Gigabit Ethernet interfaces. PR1479733
- On the ACX6360 router, the disk usage might keep increasing. PR1480217
- Memory utilization enhancement is needed. PR1481151
- Traffic loss might be observed when an ACX5448 router learned scaled with LDP-signaled prefixes. PR1482529
- ACX AUTHD process memory usage enhancement is needed. PR1482598
- BFD over Layer 2 VPN or Layer 2 circuit does not work because of the SDK upgrade to version 6.5.16. PR1483014
- On the ACX5048 router, traffic loss is observed during the unified ISSU. PR1483959
- On the ACX5448 router, the fpc process might crash. PR1485315
- On the ACX5448 routers, the Layer 2 VPN with the ethernet-circuit cross-connect input-vlan-map/output-vlan-map interface silently discards the traffic. PR1485444
- The queue statistics are not as expected after configuring the physical interface and logical interface shaping with the transmit rate and scheduler map. PR1488935
- On the ACX5448 chassis, mac-address and label mac-address might not match. PR1489034
- VPLS flood groups result in IPv4 traffic drop after the core interface flaps. PR1491261
- On the ACX5048 and ACX5096 routers, the LACP control packets might be dropped due to high CPU utilization. PR1493518
- The LSP might not come up in an externally provisioned LSP scenario. PR1494210

- There might be link issues if optic-supported speed and configured speed do not match on an interface. PR1494600
- Outbound SSH connection flap or memory leak issue might be observed during a push configuration to an ephemeral database with a high rate. PR1497575
- On the ACX5448 router, the EXP rewrite for the Layer 3 VPN sends all traffic with incorrect EXP. PR1500928
- On the ACX500 router, the SFW sessions might not get updated on the ms interfaces. PR1505089
- The PIC slot might shut down in less than 240 seconds due to the overtemperature start time being handled incorrectly. PR1506938
- The BFD session flaps with the following error message after a random time interval: **ACX_OAM_CFG_FAILED: ACX Error (oam):dnx_bfd_l3_egress_create : Unable to create egress object**. PR1513644
- The loopback filter cannot take more than 2 TCAM slices. PR1513998
- The VM process generates a core file while running the stability test in a multidimensional scenario. PR1515835
- The l2ald process crashes during the stability test with traffic on a scaled setup. PR1517074
- Tagged traffic matching the vlan-list configuration in the vlan-circuit cross-connect logical interface gets dropped in the ingress interface. PR1519568
- The **show class-of-service interface** command output does not display classifier information. PR1522941
- On the ACX5448 and ACX710 routers, the **vlan-id-list** statement might not work as expected. PR1527085
- The **show class-of-service routing-instance** does not show the configured classifier. PR1531413
- Memory leak is observed in the Local OutLif in the VPLS and CCC topology. PR1532995
- On the ACX5448 routers, the Packet Forwarding Engine crashes on the **show pfe ifd** vty command. PR1537619
- The following syslog error message is observed: **ACX_DFW_CFG_FAILED**. PR1490940
- The following error message might be observed during MPLS route add, change, or delete operation: **mpls_extra NULL**. PR1502385
- The tcpdump process generates core file after initiating the **monitor traffic interface** command. PR1485465
- Bind does not sufficiently limit the number of fetches during the referrals processing. PR1512212
- On the ACX5448 and ACX710 routers, VPLS traffic loss might be observed. PR1527231

Interfaces and Chassis

- The FPC process might crash in inline mode with CFM configured. PR1500048

MPLS

- The BGP session flaps between two directly connected BGP peers because of the wrong TCP-MSS in use. PR1493431

Routing Protocols

- The BGP route target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743
- The rpd process might report 100 percent CPU usage with BGP route damping enabled. PR1514635

VPNs

- The Layer 2 circuit neighbor might become nonresponsive in the **Ready** state at one end of the MG-LAG peer. PR1498040
- The rpd process might crash in certain conditions after deleting the Layer 2 circuit configuration. PR1502003

Resolved Issues: 20.1R1

General Routing

- On ACX5000 routers, the internal error **MacDrainTimeOut and bcm_port_update failed** is seen. PR1284590
- High CPU utilization is seen for fxpc processes with CoS changes on the aggregated Ethernet interfaces. PR1407098
- The optics module comes with Tx enabled by default. As the port is administratively disabled, the port is stopped but as the port has not been started, it does not disable Tx. PR1411015
- The l2cpd process might crash and generate a core file when interfaces flap. PR1431355
- On the ACX5448 router, the DHCP packets are not transparent over Layer 2 circuit. PR1439518
- On the ACX5448, the flexible VLAN tagging encapsulation is not supported with MPLS family, need to provide commit ERROR. PR1445046
- Fans on an ACX5448-M might not be running at the correct speed. PR1448884
- The operating state for et- interfaces does not transit from init to normal. PR1449937
- ACX5448-D interfaces support: After the 100-Gigabit and 40-Gigabit Ethernet interface are disabled, the laser output power in the output of the **show interfaces diagnostics optics** command shows some values. PR1452323
- ACX5048 SNMP polling stops after the link is flapped or the SFP transceiver is replaced and **ACX_COS_HALP(acx_cos_gport_sched_set_strict_priority:987): Failed to detach** logs will be seen. PR1455722

- ACX5448-D and ACX5448-M devices do not display airflow information and temperature sensors as expected. PR1456593
- ACX5448 Layer 2 VPN with the **encapsulation-type ethernet** configuration, stops passing traffic after a random port is added with VLAN configuration. PR1456624
- The rpd might crash if a BGP route is resolved over the same prefix protocol next hop in the inet.3 table that has both RSVP and LDP routes. PR1458595
- Route resolution is not happening when the packet size is 10,000. PR1458744
- The traffic might be discarded silently during link recovery in an open Ethernet access ring with ERPS configured. PR1459446
- On the ACX5000 router, the SNMP MIB walk for jnxOperatingTemp does not return anything for FPC in the new versions. PR1460391
- On the ACX5448-M interfaces and optics, when you enable local loopback, the 10-Gigabit Ethernet interface goes down. PR1460715
- On the ACX5448-D interfaces and optics, sometimes when you bring up the aggregated Ethernet interface, there are ARP resolution issues. PR1461485
- On ACX Series platform, the LLDP neighbor is not up on the LAG after software upgrade to Junos OS Release 18.2R3-S1. PR1461831
- Not able to add more than 16 links in a LAG. PR1463253
- Memory leak on l2cpd process might lead to l2cpd crash. PR1469635
- RED drops are seen on interfaces even without any congestion. PR1470619
- The dcpfe core is seen when disabling or enabling MACsec through ACX6360-OR scripts. PR1479710
- ACX5448 Layer 2 VPN with interface ethernet-ccc input-vlan-map/output-vlan-map can cause trafficto be discarded silently. PR1485444

Interfaces and Chassis

- MC-AE interface might show unknown status if you add the subinterface as part of the VLAN on the peer MC-AE node. PR1479012

Layer 2 Ethernet Services

- Member links state might be asychronized on a connection between the PE and CE devices in an EVPN A/A scenario. PR1463791

Routing Protocols

- The rpd might crash continuously because of memory corruption in the IS-IS setup. PR1455432

SEE ALSO

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 20.1R3 for the ACX Series documentation

Dynamic Host Configuration Protocol (DHCP)

- **Introducing DHCP User Guide**—Starting in Junos OS Release 20.1R1, we are introducing the DHCP User Guide for Junos OS routing, switching, and security platforms. This guide provides basic configuration details for your Junos OS device as DHCP Server, DHCP client, and DHCP relay agent.

  [See DHCP User Guide.]

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

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

SEE ALSO

- What's New | 13
- What's Changed | 17
- Known Limitations | 21
- Open Issues | 22
- Resolved Issues | 25
- Migration, Upgrade, and Downgrade Instructions | 33

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 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.

> **NOTE:** The following EX Series switches are supported in Release 20.1R3: EX2300, EX2300-C, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

What's New in Release 20.1R3

There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 20.1R3.

What's New in Release 20.1R2

There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 20.1R2.

What's New in Release 20.1R1

EVPN

- **Routing traffic between a VXLAN and a Layer 3 logical interface (EX4650 and QFX5120)**—Starting in Junos OS Release 20.1R1, EX4650 and QFX5120 switches support the routing of traffic between a Virtual Extensible LAN (VXLAN) and a Layer 3 logical interface. (You can configure the Layer 3 logical interface using the **set interfaces *interface-name* unit *logical-unit-number* family inet address *ip-address/prefix-length*** or the **set interfaces *interface-name* unit *logical-unit-number* family inet6 address *ipv6-address/prefix-length*** command.) This feature is enabled by default, so you do not need to take any action to enable it.

  > **NOTE:** By default, this feature is disabled on QFX5110 switches. To enable the feature on QFX5110 switches, you must perform the configuration described in Understanding How to Configure VXLANs and Layer 3 Logical Interfaces to Interoperate.

Interfaces and Chassis

- **Support for static link protection on aggregated interfaces (EX4650, QFX5120-32C, and QFX5120-48Y)**—Starting in Junos OS Release 20.1R1, you can enable link protection on aggregated interfaces for a specified static label-switched path (LSP). You can designate a primary and a backup physical link to support link protection. Egress traffic passes only through the designated primary link. This traffic includes transit traffic and locally generated traffic on the router. When the primary link fails, traffic is routed through the backup link.

  [See link-protection.]

Junos OS XML, API, and Scripting

- **The jcs:load-configuration template supports loading the rescue configuration (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **jcs:load-configuration** template supports the **rescue** parameter to load and commit the rescue configuration on a device. SLAX and XSLT scripts can call the **jcs:load-configuration** template with the **rescue** parameter set to **"rescue"** to replace the active configuration with the rescue configuration.

  [See Changing the Configuration Using SLAX and XSLT Scripts and jcs:load-configuration Template.]

Junos Telemetry Interface

- **MPLS and local routing sensor streaming support on JTI (EX2300, EX3400, EX4300, EX4600, and EX9200)**—Junos OS Release 20.1R1 provides MPLS constrained-path Label Switched Paths (LSPs), RSVP-Traffic Engineering (RSVP-TE) and local routing statistics using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. Streaming statistics are sent to an outside collector at configurable intervals.

  The following resource paths are supported:

- Local routing (resource path **/local-routes/**)
- MPLS constrained-path LSPs and RSVP-TE (resource path **/network-instances/network-instance/mpls/**)

To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC.

Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

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

- **JTI infrastructure support for (EX2300, EX2300-MP, and EX3400)**—Junos OS Release 20.1R1 provides Junos telemetry interface (JTI) infrastructure support for EX2300, EX2300-MP, and EX3400 switches.

Layer 2 Features

- **Q-in-Q support on redundant trunk links using LAGs with link protection (EX4300-MP switches and Virtual Chassis)**—Starting in Junos OS Release 20.1R1, Q-in-Q is supported on redundant trunk links (also called "RTGs") using LAGs with link protection. Redundant trunk links provide a simple solution for network recovery when a trunk port on a switch goes down. In that case, traffic is routed to another trunk port, keeping network convergence time to a minimum.

  Q-in-Q support on redundant trunk links on a LAG with link protection also includes support for the following items:

  - Configuration of flexible VLAN tagging on the same LAG that supports the redundant links configurations
  - Multiple redundant links configurations on one physical interface
  - Multicast convergence

  [See Q-in-Q Support on Redundant Trunk Links Using LAGs with Link Protection.]

Multicast

- **PIM with IPv6 multicast traffic (EX4650 and QFX5120-48Y)**—Starting in Junos OS Release 20.1R1, EX4650 and QFX5120-48Y switches support Protocol Independent Multicast (PIM) with IPv6 multicast traffic as follows:

  - PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM sparse-dense mode (PIM-SDM)
  - PIM any-source multicast (PIM-ASM) and PIM source-specific multicast (PIM-SSM)
  - Static, embedded, and anycast rendezvous points (RPs)

[See PIM Overview.]

Routing Policy and Firewall Filters

- **Support for flexible-match-mask match condition (EX4650 and QFX-Series)**—Starting with Junos OS Release 20.1R1, for EX4650, QFX5120-32C, and QFX5120-48Y switches, the **flexible-match-mask** match condition in firewall filters is supported for the **inet**, **inet6**, and **ethernet-switching** families. With this feature, you can configure a filter by specifying the length of the match (4 bytes maximum) starting from a Layer 2 or Layer 3 packet offset.

  [See Firewall Filter Flexible Match Conditions.]

Storage and Fibre Channel

- **FIP snooping (EX4650-48Y and QFX5120-48Y)**—Starting in Junos OS Release 20.1R1, EX4650-48Y and QFX5120-48Y switches support Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) snooping. With FIP snooping enabled on these switches, you prevent unauthorized access and data transmission to a Fibre Channel (FC) network by permitting only those servers that have logged in to the FC network to access the network. You enable FIP snooping on FCoE VLANs when the switch is being used as an FCoE transit switch that connects FC initiators (servers) on the Ethernet network to FCoE forwarders at the FC storage area network (SAN) edge.

  [See Understanding FCoE Transit Switch Functionality and Understanding VN_Port to VN_Port FIP Snooping on an FCoE Transit Switch.]

System Management

- **Change status LED for network port to chassis beacon light (EX4300-48MP switch and EX4300-48MP Virtual Chassis)**—By default, when a network port and its associated link are active, the status LED for that port blinks green 8 times per second. Starting in Junos OS Release 20.1R1, you can use the **request chassis beacon** command to slow down the current blinking rate to 2 blinks per second. The slower-blinking and steadier green light acts as a beacon that leads you to an EX4300-48MP switch or a particular port in a busy lab.

  Using options with the **request chassis beacon** command, you can do the following for one or all network port status LEDs on a specified FPC:

  - Turn on the beacon light for:
    - 5 minutes (default)
    - A specified number of minutes (1 through 120)
  - Turn off the beacon light:
    - Immediately
    - After a specified number of minutes (1 through 120)

  After the beacon light is turned off, the blinking rate for the network port's status LED returns to 8 blinks per second.

[See request chassis beacon.]

Virtual Chassis

- **Virtual Chassis support for up to four member switches (EX4650)**—Starting in Junos OS Release 20.1R1, you can interconnect up to four EX4650-48Y switches into a Virtual Chassis managed as a single device. The Virtual Chassis:
  - Contains only EX4650-48Y switches.
  - Has two member switches in Routing Engine role (master, backup) and the remaining members in linecard role.
  - Supports 100GbE QSFP28 or 40GbE QSFP+ ports on the front panel (ports 48 through 55) as Virtual Chassis ports (VCPs).
  - Supports NSSU.

  A EX4650-48Y Virtual Chassis with two to four members now also supports the following protocol features that were not previously supported on a two-member EX4650-48Y Virtual Chassis:
  - IEEE 802.1X authentication
  - Layer 2 port security features, including IP source guard, IPv6 router advertisement (RA) guard, DHCP, and DHCP snooping
  - MPLS
  - Redundant trunk groups (RTG)

  EX4650-48Y Virtual Chassis has limitations on protocol feature support compared to the standalone switch. The following protocol features are not supported:
  - EVPN-VXLAN
  - Junos telemetry interface (JTI)
  - Multichassis link aggregation (MC-LAG)
  - Priority-based flow control (PFC)

  Configuration and operation are the same as for other EX Series and QFX Series Virtual Chassis.

  [See Virtual Chassis Overview for Switches, 802.1X Authentication, MPLS Overview, DHCP Snooping, Understanding DHCP Snooping (ELS), Understanding IP Source Guard for Port Security on Switches, and Understanding IPv6 Router Advertisement Guard.]

SEE ALSO

What's Changed | 40

Known Limitations | 46

What's Changed

IN THIS SECTION

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

What's Changed in Release 20.1R3

EVPN

- **IGMP snooping options has changed hierarchy level**—Junos OS has moved the following options from the **edit 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

- **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

- **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.]

- **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).]

MPLS

- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the cspf-backoff-time statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

  [See cspf-backoff-time.]

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..]

- **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.]

Routing Protocols

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the lsdist.0 and lsdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, we added multiple secondary loopback addresses in the traffic engineering database to the lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router ID.

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.]

What's Changed in Release 20.1R2

General Routing

- Loading of the default configurations in a RIFT package causes the following changes:

  1. Output of the **show rift node status** command displays the node ID in hexadecimal number even though the node ID is configured in decimal, hexadecimal, or octal number.

  2. Some of the DDoS default configurations change because of the DDoS protection interferes with the RIFT BFD operation.

- **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[sid='*sid-value*']/** 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 Enabling Export of Subscriber Statistics and Queue Statistics for Dynamic Interfaces and Interface-Sets.]

High Availability (HA) and Resiliency

- **IPv6 address in the prefix TIEs displayed correctly**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.
- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**— You can set the verbosity of the trace log to only show error messages using the error option at the **edit system services extension-service traceoptions level** hierarchy.

  See traceoptions (Services).

Junos Telemetry Interface (JTI)

- **LLDP ON_CHANGE statistics support with JTI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Enhanced telemetry ON_CHANGE event support provides the following LLDP attributes:

  - When LLDP is enabled on interfaces, LLDP interface counters are notified along with other interface-level attributes.

  - ON_CHANGE event reports LLDP neighbor age and custom TLVs, as well as when a neighbor is initially discovered

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

Subscriber Management and Services

- **Command to view summary information for resource monitor (MX Series routers and EX9200 line of switches)**—You can use the **show system resource-monitor** command to view statistics about the use of memory resources for all line cards or for a specific line card in the device. The command also displays information about the status of load throttling, which manages how much memory is used before the device acts to reduce consumption.

  [See show system resource-monitor and Resource Monitoring for Subscriber Management and Services.]

What's Changed in Release 20.1R1

Interfaces and Chassis

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

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**— You can set the verbosity of the trace log to only show error messages using the error option at the **edit system services extension-service traceoptions level** hierarchy.

  [**See** traceoptions (Services).]

Multicast

- **Multicast Layer 2 transit traffic statistics by multicast source and group (EX4600, EX4650, and the QFX5000 line of switches)**—Starting in Junos OS Release 20.1R1, EX4600, EX4650, and the QFX5000 line of switches provide statistics on the packet count for each multicast group and source when passing multicast transit traffic at Layer 2 with IGMP snooping. Run the **show multicast snooping route extensive** CLI command to see this count in the **Statistics: ... *n* packets** output field. The other statistics in that output field, **kBps** and **pps**, are not available (values displayed there are not valid statistics for multicast traffic at Layer 2). In earlier Junos OS releases, all three values in the **Statistics** output field for **kBps**, **pps**, and **packets** do not provide valid statistics for multicast traffic at Layer 2.

  [See show multicast snooping route.]

SEE ALSO

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for EX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- The following error message might appear: **Failed to complete DFE tuning.** This error message has no functional impact and can be ignored. PR1473280
- In a Q-in-Q environment, if xSTP is enabled on an interface having 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 and all others will be in discarding state. This might lead to traffic drop. PR1532992

EVPN

- On the EX4650 device, inter-VNI multicast is not supported in the EVPN-VXLAN edge routing model. PR1517082

Infrastructure

- File system panic might occur after repeated power loss. PR1444941
- On EX-4300MP switches, 9000 IPv6 MC routes can be installed. If you try to add more IPv6 MC routes, error messages are seen. PR1493671

Platform and Infrastructure

- On the EX4300-MP device, ge and mge ports have different color contrasts due to different vendors. PR1470312

SEE ALSO

What's New | **35**

What's Changed | **40**

Open Issues | **48**

Resolved Issues | **51**

Documentation Updates | **63**

Migration, Upgrade, and Downgrade Instructions | **64**

Open Issues

IN THIS SECTION

Learn about open issues in this release for EX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

Infrastructure

- On EX Series switches, If you configure a large-scale number of firewall filters on some interfaces, the FPC might crash and generate core files. PR1434927
- On EX 9251 switches, **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
- On EX4300-MP switches, 9000 IPv6 MC routes can be installed. If you try to add more IPv6 MC routes, error messages are seen. PR1493671
- A double free vulnerability in the software forwarding interface daemon (sfid) process allows an adjacently-connected attacker to cause a Denial of Service (DoS) by sending a crafted ARP packet to the device. PR1497768
- The following error message is observed while loading the kernel: **GEOM: mmcsd0s.enh: corrupt or invalid GPT detected**. PR1549754
- VLAN translation (vlan mapping) does not work for CFM (0x8902) and EAPOL (0x888e). PR1580129

Interfaces and Chassis

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

Junos Fusion Provider Edge

- On Junos fusion system, intermediate traffic drop is sometimes seen between AD and SD when sFlow is enabled on the ingress interface. When sFlow technology is enabled, the original packet is getting corrupted for those packets that hit the sFlow filter This is due to few packets transmitted from the egress of AD1 is short of FCS (4 bytes) + 2 bytes of data, this leads to the drop of the packets. It is seen that the normal data packets are of size 128 bytes while the corrupted packet is 122 bytes. PR1450373

Layer 2 Features

- GARPs were being sent whenever there was a MAC (fdb) operation (add or delete). This is now updated to send GARP when the interface is up and Layer 3 interface is attached to the VLAN. PR1192520
- On EX series with third party chip based Packet Forwarding Engine, if MC-LAG is configured, and the ICL interface is a physical interface instead of an aggregated Ethernet interface, after one of the child links in multichassis aggregated Ethernet (MC-AE) interface on one of MC-LAG peers is disabled, the MAC addresses learnt from MC-LAG client device might keep flapping between the ICL interface and MC-AE interface. It could cause traffic drop when MAC addresses are learnt on ICL interface. PR1582473

Layer 2 Ethernet Services

- If the **forward-only** is set within dhcp-reply in a device as a DHCP relay agent, the DHCP DECLINE packets that are broadcasted from the DHCP client are dropped and not forwarded to the DHCP server. PR1429456

Platform and Infrastructure

- In a message queuing telemetry transport (MQTT) scenario, the memory leakage (about 4000 memory leakage every 30 seconds) might be seen. However, on long run, this uses high memory which can indirectly impact other daemons running. PR1324531
- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. PR1362609
- On an EX9208 switch, 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
- Unicast RPF in strict mode or ICMP redirect does not work properly. PR1417546

- A minimal traffic loss of ~100 pps is seen on EX9208 switches when the packets are sent between FPCs. This is due to random drops happening in the fabric. Amount of drop varies on the line rate and occurs less frequently. PR1429714
- On EX9214 switches, if the MACsec-enabled link flaps after reboot, the following error message is observed: **errorlib_set_error_log(): err_id(-1718026239)**. PR1448368
- On EX9208 switches, 33 percent degradation in MAC learning rate is seen in Junos OS Release 19.3R1 onwards while comparing with Junos OS Release 18.4R1. PR1450729
- In overall commit time, the evaluation of mustd constraints is taking two seconds more than usual. This is because the **persist-group-inheritance** feature has been made as a default feature. Eventually, this feature helps improve the subsequent commit times for scaled configurations significantly. The **persist-group-inheritance** feature is useful in customer scenarios where groups and nested groups are used extensively. In those scenarios, the group inheritance paths are not built every time, thus subsequent commits are faster. PR1457939
- On EX4300 switches, when packets entering a port exceed a size of 144 bytes, they might get dropped in very few cases. PR1464365
- The following message may be seen in chassisd log after rebooting or configuration changing, and so on: **re_tvp_builtin_fwinfo_update: Unable to get firmware version**. PR1471938
- The following syslog might be observed: **Failed to complete DFE tuning** . This message has no functional impact and can be ignored. PR1473280
- Classifiers binding applied on wildcard gets overwritten by a different classifier type when applied on a single interface. PR1490699
- While verifying Last-change op-state value through XML, rpc-reply message is inappropriate. PR1492449
- SNMP POE MIB walk produce withers no results or sometimes result from the master Virtual Chassis whenever the Virtual Chassis is renamed as one. PR1503985
- On the EX4300-48MP device, the reboot time, FPC uptime, and interface uptime are degraded by 20 percent when compared with Junos OS Releases 19.1R3, 19.2R2, and 19.4R2. PR1514364
- The MAC addresses might fail aging out under a Virtual Chassis environment where a large number of MAC addresses are learned. This issue was observed with MAC entries 280,000 in the Virtual Chassis devices. PR1558128
- EX2300 switches show high FPC CPU usage, however the system processes and kernel CPU usage does not add up to the overall FPC usage. This is a cosmetic issue with calculation of FPC CPU usage that has been resolved in newer releases of Junos OS Release 21.1R1 and later. PR1567438
- Observing traffic drop during unified ISSU due to LAG interface flap. PR1569578
- FPC core is generated at **dfw_term_cc_list_loop_init , dfw_term_cc_detect_loop , dfw_term_filter_process**. This issue might be seen only in back to back GRES in about more than 40 to 50 iterations. No workaround available and FPC gets restarted. PR1579182

Virtual Chassis

- On EX4300 Virtual Chassis platform, the virtual chassis ports might go down after the image upgrade. This issue is seen in a scenario when QSFP+-40G-SR4, QSFP+-40G-LR4, or QSFP+40GE-LX4 is used as VCP. PR1579430

SEE ALSO

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in Junos OS main and maintenance releases for EX 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: 20.1R3

Forwarding and Sampling

- The configuration archive transfer-on-commit fails. PR1563641

General Routing

- DHCP discover packets might be dropped if the DHCP inform packet is received first. PR1542400

High Availability (HA) and Resiliency

- The ksyncd process generates core files while applying the configuration to logical interfaces. PR1551777

Infrastructure

- On EX4600 and EX4300 Virtual Chassis or Virtual Chassis fabric, the VSTP configuration device goes unreachable and becomes nonresponsive after commit. PR1520351
- On 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

Interfaces and Chassis

- MC-AE interfaces might go down if same VRRP group-id is configured on multiple IRB units. PR1575779

Layer 2 Ethernet Services

- OSPF and OSPFv3 adjacency uptime is more than expected after NSSU upgrade and outage is higher than the expected. PR1551925

Platform and Infrastructure

- On EX Series platforms using chipset with SFP+ implemented, interface on the platforms might be in active status when TX or RX connector is removed. As a result, traffic might get droped. PR1495564
- A master Routing Engine reconnect might be seen on EX4300-48MP platform. PR1499771
- The DHCP traffic might not be forwarded correctly when DHCP sends unicast packets. PR1512175
- Packet drops might be seen with all commit events for 1G speed configured interface. PR1524614
- Traffic loss might be observed on interfaces in a VXLAN environment. PR1524955
- On EX3400 Virtual Chassis, console access on backup Virtual Chassis member is not allowed. PR1530106
- The lldp-receive-packet-count is not getting exchanged properly in l2pt operation for LLDP after configuring protocols. PR1532721
- The LLDP neighborship with the VoIP phones cannot be established. PR1538482
- On EX3400 and EX2300 switches, the upgrade fails due to the lack of available storage. PR1539293
- FPC might not be recognized after the power cycle (hard reboot). PR1540107
- The core dump files might be seen after the GRES or reboot. PR1541752

- The JNH memory leak could be observed on MPCs or MICs. PR1542882
- The Slaac-Snoopd child process generates core file upon multiple switchovers on the Routing Engine. PR1543181
- In every software upgrade, host needs to get upgrade. PR1543890
- On EX4300-48MP switches with Linux TVP architecture and Junos OS as VM, the Junos CLI outputs do not confirm if the Junos OS and the host kernel are compatible with each other. PR1543901
- The chip on FPC line card might crash when the system reboots. PR1545455
- On EX4300 switches, FPC crash upon receipt of specific frames on an interface without L2PT or dot1x configured. PR1545530
- FPC might not boot-up on EX9214 switches in certain conditions. PR1545838
- Receipt of specific DHCPv6 packet might cause jdhcpd process to crash and restart. PR1546166
- Classifier is not programmed in the hardware and error logs might be seen in syslog. PR1548159
- 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 **show pfe route summary hw** command shows random high free and **Used** column for IPv6 LPM (< 64)'routes. PR1552623
- The **action-shutdown** statement of storm control does not work for ARP broadcast packets. PR1552815
- The **targeted-broadcast** feature might send out duplicate packets. PR1553070
- Traffic might be dropped when a firewall filter rule uses **then vlan** as the action. PR1556198
- On EX4300 switches, script fails while committing the IPSec authentication configuration as the algorithm statement is missing. PR1557216
- The tunable optics SFP+-10G-T-DWDM-ZR does not work on EX devices. PR1561181
- On EX3400 Virtual Chassis, SMARTD pollutes syslog every 5 secs after upgrade or system reboot. PR1562396
- On EX3400VC switches, the DAEMON-7-PVIDB throws syslog messages for every 12 to 14 minutes after you upgrade. PR1563192
- On EX4650 switches, **storm control** with IRB interface might not work correctly. PR1564020
- The **Last flapped** timestamp for interface fxp0 gets reset every time when **monitor traffic interface fxp0** is executed. PR1564323
- The following internal comment is displayed: **Placeholder for QFX platform configuration**. PR1567037
- PFEX might crash when soft error recovery feature is enabled on Packet Forwarding Engine. PR1567515
- On all EX9200 switches 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 the EVPN-VXLAN routing instance. When the ASIC's next hop memory partition exhausted, the FPC might reboot. PR1571439

- DHCP packets with source IP as link-local address are dropped on EX4300 switches. PR1576022
- The LLDP packet might loss on the EX-4300MP platform if LLDP is configured on the management interface. PR1591387

Routing Protocols

- DCPFE crash might be observed while updating VRF for multicast routes during IRB uninit. PR1546745
- Sending multicast traffic to downstream receiver on MX Series-based Virtual Chassis platforms might fail. PR1555518
- The untagged packets might not work on EX Series platforms. PR1568533

Virtual Chassis

- On EX4600 and EX4300 mixed Virtual Chassis : Error message **ex_bcm_pic_eth_uint8_set** is seen when changing configuration related to interface. PR1573173

Resolved Issues: 20.1R2

Authentication and Access Control

- The authd process might have memory leak in 802.1x scenario with the RADIUS authentication. PR1503117
- On the EX2300-48MP device, the client does not receive the captive-portal success page by downloading the ACL parameter, because of the authentication failure issue. PR1504818
- The **DOT1XD_AUTH_SESSION_DELETED** event is not triggered with a single supplicant mode. PR1512724
- The 802.1x client does not go to the **Held** state when the authenticated P-VLAN is deleted. PR1516341

EVPN

- The ESI of IRB interfaces does not get updated after the autonomous-system number changes if the interface is down. PR1482790
- The l2ald memory leakage might be observed in any EVPN scenario. PR1498023
- The VXLAN function might be broken due to a timing issue. PR1502357
- Unable to create a new VTEP interface. PR1520078

General Routing

- The Virtual Chassis splits after the network topology changes. PR1427075
- The MAC pause frames keep incrementing in the receive direction if half-duplex mode on 10-Mbps or 100-Mbps speed is configured. PR1452209

- The FPC process might get disconnected from the EX3400 Virtual Chassis briefly after rebooting or upgrading. PR1467707
- On the EX4600 device, traffic loss might be seen with framing errors or runts if MACsec is configured. PR1469663
- On the EX4600 device, the DSCP marking might not work as expected if the fixed classifiers are applied to interfaces. PR1472771
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. PR1473610
- On the EX4300 device, the output of the **show security macsec statisitics** command displays high values incorrectly. PR1476719
- On the EX2300 device, the SNMP traps are not generated when the MAC addresses limit threshold is reached. PR1482709
- Incorrect frame length of 132 bytes might be captured in the packet header. PR1487876
- DHCP binding might fail when the P-VLAN is configured with a firewall to block or allow certain IPv4 packets. PR1490689
- On the EX2300 device, high CPU load due to the receipt of specific multicast packets on Layer 2 interface is observed. PR1491905
- On the EX4300 device, traffic loss might be observed in a mixed Virtual Chassis setup. PR1493258
- On the EX4650 device, traffic loss might be seen in an MC-LAG scenario. PR1494507
- The authentication session might be terminated if the PEAP request is retransmitted by an authenticator. PR1494712
- The fxpc process might crash when renumbering the primary member ID value of the EX2300 or EX3400 Virtual Chassis. PR1497523
- Outbound SSH connection flap or memory leak issue might be observed during a push configuration to an ephemeral database with a high rate. PR1497575
- Traffic might get dropped if the aggregated Ethernet member interface is deleted or added, or an SFP transceiver of the aggregated Ethernet member interface is unplugged or plugin. PR1497993
- In some cases, if we have an OSPF session on the IRB over LAG interface with a 40-Gigabit Ethernet port as member, the session gets stuck at restart. PR1498903
- On the EX4300, EX3400, and EX2300 Virtual Chassis with NSB and xSTP enabled, continuous traffic loss might be observed while performing GRES. PR1500783
- The mge interface might still stay up while the far end of the link goes down. PR1502467
- LLDP is not acquired when the native VLAN ID and the tagged VLAN ID are the same on a port. PR1504354
- The isolated VLAN from RADIUS is not deleted when the interface flaps. PR1506427
- The output VLAN push might not work. PR1510629

- LLDP might not work when P-VLAN is configured. PR1511073
- On the EX4300 device, LACP goes down after a Routing Engine switchover if MACsec is enabled on the LAG members. PR1513319
- The 100-Mbps SFP-FX transceiver is not supported on a satellite device in the Junos fusion setup. PR1514146
- 802.1x memory leak is observed. PR1515972
- The dcpfe process might crash due to memory leak. PR1517030
- MPPE-Send or Recv-key attribute is not extracted correctly by dot1xd. PR1522469
- Drops and dropped packets counters in the output value of the **show interface extensive** command are counted twice. PR1525373
- On the EX2300 device, the following PoE message is observed: **poe_get_dev_class: Failed to get PD class info**. PR1536408
- Traffic impact might be observed on the EVPN-VXLAN scenario due to ARP reply not working properly with **native-vlan-id** configured. PR1483167
- IRB MAC does not get programmed in hardware when the MAC persistence timer expires. PR1484440
- BIND does not sufficiently limit the number of fetches during the referrals processing. PR1512212
- Memory leakage is observed while processing specific DHCP packets. PR1514145
- On the EX4300-MP router, ARP learning issue might be observed when configuring the Layer 3 gateway interfaces. PR1514729

High Availability (HA) and Resiliency

- Kernel generates core file on the backup Routing Engine causing traffic drop if multicast-MAC is configured on the IRB interface. PR1467847

Infrastructure

- On the EX2300 and EX3400 devices, the kernel might generate core files when deactivating the daemon. PR1483644
- The fxpc might crash when configuring scaled configuration with 4093 VLANs. PR1493121
- On the EX4600 device, the IP communication between directly connected interfaces might fail. PR1515689
- DUT did not receive the LLDP packet from phone. PR1538482
- On the EX4600 and EX4300 Virtual Chassis or Virtual Chassis Fabric, the VSTP configurations device goes unreachable and becomes nonresponsive after commit. PR1520351

Interfaces and Chassis

- The following syslog message is observed after MX-VC local or global switchover: **scchassisd[ ]: CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments fru_send_msg Global FPC x**. PR1428254
- The MC-LAG configuration-consistency ICL-configuration might fail after committing some changes. PR1459201
- A stale IP address might be seen after a specific order of configuration changes in the logical-systems scenario. PR1477084
- Traffic might get dropped as the next hop points to the ICL even though the local MC-LAG is up. PR1486919

Junos Fusion Enterprise

- The following error message is observed with duplicate ECID values for cluster or extended ports on member ports of the same cluster: **jnh_dot1br_ktree_entry_create(1098): Entry Already Exists** . PR1408947
- The SDPD generates core files at **vfpc_all_eports_deletion_complete vfpc_dampen_fpc_timer_expiry**. PR1454335

Junos Fusion Satellite Software

- On the EX4300 device, the temperature sensor alarm is seen. PR1466324

Layer 2 Ethernet Services

- Issues with the DHCPv6 relay processing confirm and reply packets are observed. PR1496220
- Default-route might not be added to the Juniper OS device configured as the DHCPv4 client device. PR1504931

Layer 2 Features

- The third VLAN tag does not get pushed onto the stack. Instead, it gets swapped. PR1469149
- Traffic imbalance might be observed if hash-params is not configured. PR1514793
- The MAC address in the hardware table might become out of synchronization between the primary devices and the member devices in the Virtual Chassis after the MAC flaps. PR1521324
- The dcpfe or the FPC process might crash due to the memory leakage during the VLAN addition or deletion operation. PR1505239

MPLS

- BGP session flaps between two directly connected BGP peers because of the wrong TCP-MSS in use. PR1493431

Platform and Infrastructure

- MAC learning under bridge-domain stops after the MC-LAG interface flaps. PR1488251
- The traffic destined to VRRP VIP might be dropped after the IRB interface is disabled on the initial VRRP primary device. PR1491348
- IPv6 neighbor solicitation packets might be dropped in a transit device. PR1493212
- Packets get dropped when the next hop is an IRB-over-lt interface. PR1494594
- On the EX4300 device, NSSU might fail due to a storage issue on the **/var/tmp** directory. PR1494963
- On the EX4300 device, high CPU load due to receipt of specific IPv4 packets is observed. PR1495129
- On the EX4300 device, traffic loss might be seen with framing errors or runts if MACsec is configured. PR1502726
- On the EX4300 device, the redirected IP traffic is being duplicated. PR1518929
- LLDP neighborship might not come up on EX4300 non-aggregated Ethernet interfaces. PR1538401
- Memory leaks in the Packet Forwarding Engine due to the flapping of the 802.1X authenticator port interface. PR1480706
- Trio-based MPC memory leaks when the IRB interface is mapped to a VPLS instance or a Bridge-Domain. PR1525226
- On the EX4300-VC devices, the FBF functionality might be broken after rebooting the Virtual Chassis or on modifying the IRB configuration. PR1531838

Routing Protocols

- The MUX state in the LACP interface does not go to the **Collecting** and **Distributing** states and remains in the **Attached** state after enabling the aggregated Ethernet interface. PR1484523
- The FPC process goes to the **NotPresent** state after upgrading the Virtual Chassis or Virtual Chassis Fabric. PR1485612
- The BGP route target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743
- On the EX4300-MP and EX4600 devices, high CPU load due to receipt of specific Layer 2 frames in EVPN-VXLAN deployment. PR1495890
- Firewall filter does not work in certain conditions under the Virtual Chassis setup. PR1497133
- The rpd might report 100 percent CPU usage with BGP route damping enabled. PR1514635
- Packet loss might be observed while verifying traffic from access to core network for IPv4 or IPv6 interfaces. PR1520059

- The OSPFv3 adjacency should not be established when IPsec authentication is enabled. PR1525870
- Packets drop might be observed when the multicast MAC with static ARP is configured on one IRB interface. PR1489374

User Interface and Configuration

- On the EX2300 and EX3400 devices, installing J-Web application package might fail. PR1513612
- J-Web does not display the correct flow-control status. PR1520246

Virtual Chassis

- On the EX4650 device, the following error message is observed during booting: **kldload: an error occurred while loading the module**. PR1527170

Resolved Issues: 20.1R1

Authentication and Access Control

- On EX4600 and EX4300 switches, MAC entry is missing in the Ethernet switching table for Mac-radius client in server fail scenario when tagged is sent for two client. PR1462479

Class of Service (CoS)

- Shaping does not work after the reboot if **shaping-rate** is configured. PR1432078
- The traffic is placed in network-control queue on an extended port even if it comes in with different DSCP marking. PR1433252

EVPN

- The rpd might crash after the EVPN-related configuration is changed. PR1467309

Forwarding and Sampling

- Type 1 ESI/AD route might not be generated locally on the EVPN PE device in the **all-active** mode. PR1464778

General Routing

- The l2cpd process might crash and generate a core file when interfaces flap. PR1431355
- MicroBFD flap is seen when a QSFP transceiver is inserted into other port. PR1435221
- EX4600 Virtual Chassis does not come up after the Virtual Chassis port fiber connection is replaced with a DAC cable. PR1440062
- MAC addresses learned on an RTG might not be aged out after a Virtual Chassis member reboots. PR1440574
- Except one aggregated Ethernet member link, the other links do not send out sFlow sample packets for ingress traffic. PR1449568

- On EX3400 switches with half-duplex mode on 10-Mbps or 100-Mbps speed at medium traffic egress, traffic flow might stop on the port and MAC pause frames will be incrementing in the receive direction. PR1452209
- The l2ald and eventd processes are hogging 100 percent after the **clear ethernet-switching table** command is issued. PR1452738
- A firewall filter might not be applied in a particular Virtual Chassis or Virtual Chassis Fabric member as TCAM is running out of space. PR1455177
- Packet drop might be seen after removing and reinserting the SFP transceiver of the 40G uplink module ports. PR1456039
- Link-up delay and traffic drop might be seen on mixed SP L2/L3 and EP L2 type configurations. PR1456336
- The **syslog timeout connecting to peer database-replication** message is generated when the **show version detail** command is issued. PR1457284
- Overtemperature SNMP trap messages appear after an update even though the temperature is within the system thresholds. PR1457456
- The correct VoIP VLAN information in LLDP-MED packets might not be sent after commit if dynamic VoIP VLAN assignment is used. PR1458559
- The FXPC process might crash due to several BGP IPv6 session flaps. PR1459759
- On EX2300 and EX3400 switches, storage space limitation leads to image installation failure during phone home. PR1460087
- MAC addresses learned on redundant trunk group (RTG) might not be aged out after the aging time if the source interface is configured as RTG. PR1461293
- RTG link is down for nearly 20 seconds when the backup node is rebooting. PR1461554
- Configuring any combination of VLANs and interfaces under VSTP/MSTP might cause the VSTP/MSTP-related configuration to fail. PR1463251
- The Virtual Chassis function might be broken after an upgrade on EX2300 and EX3400 devices. PR1463635
- A few command lines to disable MAC learning are not working. PR1464797
- The jdhcpd might consume a high CPU and no further subscribers can be brought up if there are more than 4000 DHCP relay clients in the MAC move scenario. PR1465277
- On EX2300 switches, an FXPC core file is seen after mastership election based on the user's priority. PR1465526
- The broadcast and multicast traffic might be dropped over an IRB or a LAG interface in a Virtual Chassis scenario. PR1466423
- The MAC move message might have an incorrect **from** interface when MAC moves rapidly. PR1467459

- Optics measurements might not be streamed for interfaces of a PIC over JTI. PR1468435
- SSH session closes while you check the **show configuration | display set** command for both local and non-local users. PR1470695
- EX3400 switch is advertising only 100 Mbps when a speed of 100 Mbps is configured with autonegotiation enabled. PR1471931
- On EX4600 switches, the shaping of CoS does not work after reboot. PR1472223
- On EX3400 switches, CoS 802.1p bits rewrite might not happen in Q-in-Q mode. PR1472350
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. PR1473685
- The dhcpd process might crash in a Junos fusion environment. PR1478375
- MX Series with MPCs/MICs based line-card might crash when there is a bulk route update failure in a corner case. PR1478392
- TFTP installation from loader prompt might not succeed on EX Series devices. PR1480348
- In an EVPN-VXLAN scenario, ARP request packets for an unknown host might be dropped in remote PE device. PR1480776

Infrastructure

- EX2300 switches might stop forwarding traffic or responding to the console. PR1442376
- On EX4300 switches, the CLI configuration **set chassis routing-engine on-disk-failure disk-failure-action (reboot | halt)** is not supported. PR1450093
- EX Series switches might not come up properly after reboot. PR1454950
- On EX4600 and EX4300 Virtual Chassis, error messages related to soft reset of port due to queue buffers being stuck could be seen. PR1462106
- Traffic is dropped on an EX4300-48MP device acting as a leaf device in a Layer 2 IP fabric EVPN-VXLAN environment. PR1463318
- EX3400 switches might reboot because of lack of watchdog patting. PR1469400
- In an EX2300 Virtual Chassis scenario, continuous dcpfe error messages and eventd process hog might be seen. PR1474808

Interfaces and Chassis

- VRRPv6 state is flapping with init and idle states after configuring **vlan-tagging**. PR1445370
- Traffic might be forwarded to incorrect interfaces in an MC-LAG scenario. PR1465077
- Executing commit might become unresponsive due to stuck device control process. PR1470622

Junos Fusion Enterprise

- Loop detection might not work on extended ports in Junos fusion scenarios. PR1460209

Junos Fusion Satellite Software

- In Junos fusion for enterprise, the dpd crash might be observed on satellite devices running SNOS. PR1460607

Layer 2 Features

- MAC or ARP learning might not work for copper base SFP-T transceivers on EX4600 switches. PR1437577
- The Link Layer Discovery Protocol (LLDP) function might fail when a Juniper device connects to a non-Juniper device. PR1462171
- After rebooting, an FXPC core file might be seen when committing the configuration. PR1467763
- Traffic might be affected if composite next-hop is enabled. PR1474142

Layer 2 Ethernet Services

- Member links state might be asynchronized on a connection between PE and CE devices in an EVPN A/A scenario. PR1463791

Platform and Infrastructure

- NSSU causes traffic loss again after the backup to master transitions. PR1448607
- In a Virtual Chassis scenario, the IRB traffic might get dropped after master switchover. PR1453025
- The OSPF neighbor might go down when mDNS/PTP traffic is received at a rate higher than 1400 pps. PR1459210
- ERP might not revert to IDLE state after reload or reboot of multiple switches. PR1461434
- On EX4300 Virtual Chassis, traffic loss might be observed longer than 20 seconds when performing NSSU. PR1461983
- On EX2300 and EX3400 switches, the upgrade might fail as there is not enough space. PR1464808
- On EX4300 switches, IGMP reports are dropped when mixed enterprise and service provider configuration styles are used. PR1466075
- On EX4300 switches, an input firewall filter attached to isolated or community VLANs fails to match dot1p bits on the VLAN header. PR1478240
- Virtual Chassis VRRP peer drops packets destined to the VRRP VIP after IRB is disabled. PR1491348

Routing Protocols

- Host-destined packets with the filter log action might not reach the Routing Engine if log or syslog is enabled. PR1379718
- On EX9208 platforms, BGP IPv4 or IPv6 convergence and RIB install or delete time are degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. PR1414121

- The **other querier present interval** timer cannot be changed in an IGMP/MLD snooping scenario. PR1461590

User Interface and Configuration

- Problem with access to J-Web after updating from Junos OS Release 18.2R2 to Release 18.2R3. PR1454150
- Error message **umount: unmount of /.mount/var/val/chroot/packages/mnt/jweb-ex32-d2cf6f6b failed: Device busy** is seen when Junos OS is upgraded with the **validate** option. PR1478291

SEE ALSO

What's New | 35

What's Changed | 40

Known Limitations | 46

Open Issues | 48

Documentation Updates | 63

Migration, Upgrade, and Downgrade Instructions | 64

Documentation Updates

IN THIS SECTION

- Dynamic Host Configuration Protocol (DHCP) | 64

This section lists the errata and changes in Junos OS Release 20.1R3 documentation for the EX Series.

Dynamic Host Configuration Protocol (DHCP)

- **Introducing DHCP User Guide**—Starting in Junos OS Release 20.1R1, we are introducing the DHCP User Guide for Junos OS routing, switching, and security platforms. This guide provides basic configuration details for your Junos OS device as DHCP Server, DHCP client, and DHCP relay agent.

  [See DHCP User Guide.]

SEE ALSO

Resolved Issues | 51

What's Changed | 40

Known Limitations | 46

Open Issues | 48

Resolved Issues | 51

Migration, Upgrade, and Downgrade Instructions | 64

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for the EX Series. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network. For information about software installation and upgrade, see the Installation and Upgrade Guide.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

SEE ALSO

- What's New | 35
- What's Changed | 40
- Known Limitations | 46
- Open Issues | 48
- Resolved Issues | 51
- Documentation Updates | 63

Junos OS Release Notes for JRR Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 for JRR Series. They describe new and changed features, limitations, and known and resolved problems in the hardware.

You can 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 for JRR Series in Junos OS Release 20.1R3.

SEE ALSO

- What's Changed | 66
- Known Limitations | 67
- Open Issues | 67
- Resolved Issues | 68
- Documentation Updates | 69
- Migration, Upgrade, and Downgrade Instructions | 69

What's Changed

There are no changes in behavior and syntax for JRR Series in Junos OS Release 20.1R3.

SEE ALSO

- What's New | 66
- Known Limitations | 67
- Open Issues | 67
- Resolved Issues | 68
- Documentation Updates | 69
- Migration, Upgrade, and Downgrade Instructions | 69

Known Limitations

There are no known limitations for JRR Series in Junos OS Release 20.1R3.

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

SEE ALSO

- What's New | 66
- What's Changed | 66
- Open Issues | 67
- Resolved Issues | 68
- Documentation Updates | 69
- Migration, Upgrade, and Downgrade Instructions | 69

Open Issues

There are no open issues for JRR Series in Junos OS 20.1R3 Release.

SEE ALSO

- What's New | 66
- What's Changed | 66
- Known Limitations | 67
- Resolved Issues | 68
- Documentation Updates | 69
- Migration, Upgrade, and Downgrade Instructions | 69

Resolved Issues

IN THIS SECTION

Learn about resolved issues for JRR Series in Junos OS 20.1R3 Release.

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

Resolved Issues: 20.1R3

General Routing

- When performing the **request system power-off** command on JRR200 platform, the device might reboot instead of shutting down. Then, the traffic might be lost until the system reboot. Additionally, if **request system halt** command is executed, some of the physical interfaces remains up which might cause traffic impact. PR1534795
- The factory default configuration on JRR200 does not have vendor-id option configured on the interfaces (em0; em2-em9) and as a result DHCP option 60 is not sent out to the DHCP/ZTP server during the ZTP process. PR1582038

Resolved Issues: 20.1R2

General Routing

- Link state of virtual em interfaces in Junos OS might not reflect the true link status of corresponding physical interfaces in the Linux host. PR1492087
- The tcp_timer_keep log messages flood continuously on JRR200. PR1533168

Resolved Issues: 20.1R1

General Routing

- JRR200: USB install image is not working. PR1471986

SEE ALSO

- What's New | 66
- What's Changed | 66
- Known Limitations | 67
- Open Issues | 67
- Documentation Updates | 69
- Migration, Upgrade, and Downgrade Instructions | 69

Documentation Updates

There are no errata or changes in Junos OS Release 20.1R3 documentation for JRR200 Route Reflectors.

SEE ALSO

- What's New | 66
- What's Changed | 66
- Known Limitations | 67
- Open Issues | 67
- Resolved Issues | 68
- Migration, Upgrade, and Downgrade Instructions | 69

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

This section contains the upgrade and downgrade support policy for Junos OS for the 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 Installation and Upgrade Guide.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

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

SEE ALSO

- What's New | 66
- What's Changed | 66
- Known Limitations | 67
- Resolved Issues | 68
- Open Issues | 67
- Documentation Updates | 69

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

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

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

What's New

There are no new features or enhancements to existing features in Junos OS Release 20.1R3 for Junos fusion for enterprise.

> **NOTE:** For more information about the Junos fusion for enterprise features, see the Junos fusion for enterprise User Guide.

SEE ALSO

What's Changed

There are no changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 20.1R3 for Junos fusion for enterprise.

SEE ALSO

Known Limitations

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 20.1R3 for Junos fusion Enterprise.

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

Junos fusion for enterprise

- In a Junos fusion, intermediate traffic drop might be seen between the aggregation and satellite device when sFlow is enabled on the ingress interface. When sFlow is enabled, the original packet is corrupted for those packets that hit the sFlow filter. This is because the packets exiting the aggregation device are short of 4 bytes of FCS and 2 bytes of data. Normal data packets are 128 bytes (4 bytes for FCS, 14 bytes for Ethernet header, 20 bytes for IP header, and 90 bytes for data). The corrupted packets are 122 bytes (14 bytes for Ethernet header, 20 bytes for IP header, and 88 bytes for data). PR1450373

SEE ALSO

What's New | 71

What's Changed | 72

Open Issues | 73

Resolved Issues | 74

Documentation Updates | 75

Migration, Upgrade, and Downgrade Instructions | 75

Open Issues

There are no known issues in hardware and software in Junos OS Release for 20.1R3 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.

SEE ALSO

What's New | 71

What's Changed | 72

Known Limitations | 72

Resolved Issues | 74

Documentation Updates | 75

Migration, Upgrade, and Downgrade Instructions | 75

Resolved Issues

IN THIS SECTION

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

Resolved Issues: 20.1R3

There are no resolved issues in Junos Release 20.1R3 for Junos fusion for enterprise.

Resolved Issues: 20.1R2

- When using EX9200 as the aggregation device, you might see the error, "**jnh_dot1br_ktree_entry_create(1098): Entry Already Exists**," continuously. This is caused by EX2300 sattilite devices having duplicate ECID values for the cluster/extended ports across members of same cluster devices. PR1408947
- The SDPD process generates core files at **vfpc_all_eports_deletion_complete vfpc_dampen_fpc_timer_expiry**. PR1454335
- In a Junos fusion, an EX4300 acting as the satellite device is triggering the temperature sensor alarm on multiple satellite device modules connected to same aggregation device. PR1466324

Resolved Issues: 20.1R1

- Loop detection might not work on extended ports in a Junos fusion for enterprise scenario. PR1460209
- The dpd process might generate a core file on satellite devices in Junos fusion for enterprise. PR1460607

SEE ALSO

What's New | 71
What's Changed | 72
Known Limitations | 72
Open Issues | 73
What's Changed | 72
Migration, Upgrade, and Downgrade Instructions | 75

Documentation Updates

There are no errata or changes in Junos OS Release 20.1R3 for documentation for Junos fusion for enterprise.

SEE ALSO

What's New | 71
What's Changed | 72
Known Limitations | 72
Open Issues | 73
Resolved Issues | 74
Migration, Upgrade, and Downgrade Instructions | 75

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/junos-install-ex92xx-x86-64-18.3B1.n.tgz
```

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

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-18.3B1.n-limited.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

There are multiple methods to upgrade or downgrade satellite software in your Junos fusion 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

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

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

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

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

Downgrading from Junos OS

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

SEE ALSO

- What's New | 71
- What's Changed | 72
- Known Limitations | 72
- Open Issues | 73
- Resolved Issues | 74
- Documentation Updates | 75

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

Learn about new features introduced in this release for Junos fusion Provider Edge routers.

What's New in 20.1R3 Release

There are no new features or enhancements to existing features for Junos fusion for provider edge in Junos OS Release 20.1R3.

What's New in 20.1R2 Release

There are no new features or enhancements to existing features for Junos fusion for provider edge in Junos OS Release 20.1R2.

What's New in 20.1R1 Release

Hardware

- **Support for MX10008 and MX10016**—Starting in Junos OS Release 20.1R1, Junos fusion for provider edge supports the use of an MX10008 or MX10016 router as an aggregation device which acts as the single point of management for all devices in the Junos fusion.

  [See Understanding Junos Fusion Provider Edge Software and Hardware Requirements.]

SEE ALSO

What's Changed

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 Junos fusion for provider edge.

SEE ALSO

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.1R3 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.

SEE ALSO

Open Issues

IN THIS SECTION

Learn about open issues in this release 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.

Junos Fusion Provider Edge

- On a Junos fusion system, intermediate traffic drop might be seen between the aggregation and satellite device when SFlow is enabled on the ingress interface. When SFlow is enabled, the original packet is corrupted for those packets that hit the SFlow filter. This is because the packets egressing the aggregation device are short of 4 bytes of FCS 2 bytes of data. Normal data packets are 128 bytes (4 bytes for FCS, 14 bytes for Ethernet header, 20 bytes for IP header and 90 bytes for data). The corrupted packets are 122 bytes (14 bytes for Ethernet header, 20 bytes for IP header, and 88 bytes for data).PR1450373

SEE ALSO

What's New | 81

What's Changed | 82

Known Limitations | 83

Resolved Issues | 84

Documentation Updates | 85

Migration, Upgrade, and Downgrade Instructions | 86

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: Release 20.1R3

There are no resolved issues in Junos OS Release 20.1R3 for Junos fusion for provider edge.

Resolved Issues: Release 20.1R2

Junos Fusion for Provider Edge

- The statistics of the extended ports on the satellite device cluster might show wrong values from the aggregation device. PR1490101

Resolved Issues: Release 20.1R1

There are no resolved issues in Junos OS Release 20.1R1 for Junos fusion for provider edge.

SEE ALSO

- What's New | 81
- What's Changed | 82
- Known Limitations | 83
- Open Issues | 83
- Documentation Updates | 85
- Migration, Upgrade, and Downgrade Instructions | 86

Documentation Updates

There are no errata or changes in Junos OS Release 20.1R3 documentation for Junos fusion for provider edge.

SEE ALSO

- What's New | 81
- What's Changed | 82
- Known Limitations | 83
- Open Issues | 83
- Resolved Issues | 84

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 20.1R3 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/jinstall64-20.1R3.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.1R3.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/jinstall64-20.1R3.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.1R3.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 20.1R3 **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 20.1R3, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Downgrading from Junos OS Release 20.1

To downgrade from Release 20.1 to another supported release, follow the procedure for upgrading, but replace the 20.1 **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.

SEE ALSO

- What's New | **81**
- What's Changed | **82**
- Known Limitations | **83**
- Open Issues | **83**
- Resolved Issues | **84**
- Documentation Updates | **85**

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

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 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 this release for MX Series routers.

What's New in 20.1R3 Release

There are no new features or enhancements to existing features for MX Series routers in Junos OS Release 20.1R3.

What's New in 20.1R2 Release

There are no new features or enhancements to existing features for MX Series routers in Junos OS Release 20.1R2.

What's New in 20.1R1 Release

Hardware

> **NOTE:** The MX2K-MPC11E line card is supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases and in Junos OS Release 20.1R1 and later Junos OS releases. It is not supported in any Junos OS 19.4 releases.

Class Of Service

- **Hierarchical CoS support on MX2K-MPC11E line cards (MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 20.1R1, hierarchical CoS is supported on MX2K-MPC11E line cards.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Protocols and Applications Supported by the MX2K-MPC11E.]

- **Forwarding CoS (L2 classifiers, rewrite) support on MX2K-MPC11E line cards (MX2008, MX2010, and MX2020)**—Starting in Release 20.1R1, Junos OS supports forwarding CoS (L2 classifiers, rewrite) for MX Series routers with MX2K-MPC11E line cards.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Protocols and Applications Supported by the MX2K-MPC11E.]

- **Seamless MPLS CoS support for pseudowires from access node and multiservices edge (MSE) node on MX2K-MPC11E line cards (MX2008, MX2010, and MX2020)**—Starting with Junos OS Release 20.1R1, we support on the MX2K-MPC11E line card pseudowires from access node and multiservices edge

(MSE) node for MX2008, MX2010, and MX2020 routers to include seamless MPLS CoS (BA and MF classifiers, rewrite rules, schedulers, drop profiles, policers, HQoS support — interface-set, physical interface level, S-VLAN level, logical unit/C-VLAN level, and traffic-control profile).

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Protocols and Applications Supported by the MX2K-MPC11E.]

- **CoS support for forwarding class counters on MX2K-MPC11E line cards (MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 20.1R1, we support forwarding class counters on MX2K-MPC11E line cards. This feature was originally introduced in Junos OS Release 14.1.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Protocols and Applications Supported by the MX2K-MPC11E.]

- **Layer 2.5 injection of control traffic to ensure queuing on GRE tunnel with CoS settings intact (MX204 and devices installed with next-generation MPCs (MPC2E-NG and MPC3E-NG))**—Starting with Junos OS Release 20.1R1, you can configure host-injected control traffic to reache the GRE tunnel interface queues at the packet forwarding engine when the control session is over the GRE tunnel interface. This includes control protocols OSPF, BGP, PIM, RSVP, LDP, OAM, BFD, and MSDP. Injection of control traffic ensures that the kernel includes the interface ID of the GRE logical interface and the unicast next-hop ID of the corresponding GRE physical interface along with the packet that is injected into the packet forwarding engine.

  With this feature enabled, all transit packets on the GRE tunnel logical interface have the ToS copied to the outer header. To enable this feature, configure the **force-control-packets-on-transit-path** statement on the GRE tunnel logical interface.

  [See force-control-packets-on-transit-path.]

EVPN

- **Support for EVPN functionality on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS 20.1R1, you can configure MX2K-MPC11E line cards on MX2010 and MX2020 routers to support single-homed devices on an EVPN-MPLS network.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See EVPN Multihoming Overview.]

Forwarding and Sampling

- **Support for load balancing on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, the following advanced Layer 2 features are supported on MX2010 and MX2020 routers with MX2K-MPC11E line cards and Enhanced Switch Fabric Boards (SFB3s): enhanced hash-key options, consistent flow hashing, symmetrical load balancing over 802.3ad LAGs, source IP only hashing, and destination IP only hashing.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Configuring Per-Flow Load Balancing Based on Hash Values.]

General Routing

- **Support for GRE key (MX Series)**—Starting with Junos OS 20.1R1, Junos OS supports configuring a key to identify traffic flows in a GRE tunnel as defined in RFC2890. You must configure the key on the routers on both endpoints of a tunnel and create an export policy to populate the key in the forwarding table. You can configure **dynamic-tunnel-gre-key** at the **[edit routing-options dynamic-tunnels tunnel-attributes *name*]** hierarchy level.

  [See dynamic-tunnel-gre-key.]

High Availability and Resiliency

- **Unified ISSU with enhanced mode (MX240, MX480, MX960, MX2008, MX2010, MX2020)**—Starting in Junos OS Release 20.1R1, MX Series routers with MPC8E, or MPC9E line cards installed can use a new ISSU option called *enhanced mode*. Enhanced mode eliminates packet loss during the unified ISSU process by running a copy of the Junos OS software in standby mode ready to take over when software moves from an old image to a new one.

  Use the **request system software in-service-upgrade *package-name*.tgz enhanced-mode** command to use unified ISSU with enhanced mode, or the **request system software validate in-service-upgrade**

***package-name*.tgz enhanced-mode** command to verify that your device and target release are compatible with enhanced mode.

[See How to Use Unified ISSU with Enhanced Mode.]

- **Sequential upgrade for Virtual Chassis (MX240, MX480, MX960, and MX10003)**—Starting in Junos OS Release 20.1R1, MX Series Virtual Chassis configurations can use sequential upgrade to install new software releases with minimal network downtime. Sequential upgrade is an alternative to unified ISSU that installs a new release and reboots each Virtual Chassis member router one at a time. While the upgrade happens on one member router, the other member router continues to operate and handle network operations.

  To perform a sequential upgrade in an MX Series Virtual Chassis, you first issue the **request virtual-chassis upgrade protocol-backup *package-name*** command from the CLI for the Virtual Chassis master router. This initiates the upgrade process on the Virtual Chassis backup router. After the upgrade finishes on the backup router, issue the **request virtual-chassis upgrade protocol-master *package-name*** command from the backup router CLI to begin the same upgrade process for the Virtual Chassis master router.

  [See How to Use Sequential Upgrade in an MX Series Virtual Chassis.]

- **Support for BFD on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, MX2010 and MX2020 routers with the MX2K-MPC11E line card support the following BFD features:

  - Centralized BFD
  - Distributed BFD
  - Inline BFD (single-hop only)
  - Single-hop BFD
  - Multihop BFD
  - Micro BFD
  - BFD over IRB interfaces
  - BFD over pseudowire over logical tunnel and redundant logical tunnel interfaces
  - Virtual circuit connectivity verification (VCCV) BFD for Layer 2 VPNs, Layer 2 circuits, and virtual private LAN service (VPLS)

  Micro BFD at the Packet Forwarding Engine level behaves slightly differently on MX2K-MPC11E line cards. If micro BFD is enabled on an aggregated Ethernet (ae) interface, the micro BFD packets are not subjected to firewall filters for both tagged and untagged ae interfaces.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Understanding BFD for Static Routes for Faster Network Failure Detection and Understanding Distributed BFD.]

Interfaces and Chassis

- **Support for flexible tunnel interfaces (MX240, MX480, and MX960 with MPC10E; MX2010 and MX2020 with MPC11E)**—Starting in Junos OS Release 20.1R1, MX Series routers with MPC10E or MPC11E line cards support flexible tunnel interfaces (FTIs). FTIs support Layer 3 point-to-point tunnels, which use Virtual Extensible LAN (VXLAN) encapsulation with a Layer 2 pseudo header.

  To configure FTIs on your device and to enable multiple encapsulations on the FTIs, use the **vxlan-gpe** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* tunnel encapsulation]** hierarchy level.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Flexible Tunnel Interfaces Overview and vxlan-gpe (FTI).]

- **Support for ALB on multiple Packet Forwarding Engines for aggregated Ethernet bundles (MX Series MPCs)**—Starting in Junos OS Release 20.1R1, on MX Series MPCs, adaptive load balancing (ALB) for aggregated Ethernet bundles evenly redistributes the traffic load across multiple ingress Packet Forwarding Engines on the same line card, thus providing flexibility and redundancy. In earlier releases, ALB evenly redistributes traffic across all ingress traffic on a single Packet Forwarding Engine only. ALB is disabled by default.

  > **NOTE:** MPC3E does not support adaptive load balancing.

  To configure ALB, include the **adaptive** statement at the [**edit interfaces ae-interface aggregated-ether-options load-balance**] hierarchy level.

  > **NOTE:** When you configure locality bias and adaptive load balancing for aggregated Ethernet interfaces, ALB is supported per Packet Forwarding Engine and not across all Packet Forwarding Engines on the same line card. Also, you cannot revert to ALB support per Packet Forwarding Engine after you enable ALB support on multiple Packet Forwarding Engines.

  [See Configuring Adaptive Load Balancing.]

- **Adaptive load balancing on MPC10E-15C-MRATE, MPC10E-10C-MRATE, and MX2K-MPC11E line cards (MX240, MX480, MX960, and MX2020)**—Starting in Junos OS Release 20.1R1, adaptive load balancing (ALB) is supported on aggregated Ethernet bundles and ECMP links to correct traffic imbalance

among member links. ALB resolves traffic load imbalance caused by the hashing algorithm. With ALB configured on the system, traffic is balanced across member links when an imbalance is detected.

- To configure ALB on aggregated Ethernet bundles, run the **set interfaces name aggregated-ether-options load-balance adaptive tolerance** command. [See adaptive.]
- To configure ALB on ECMP links, run the **set chassis ecmp-alb tolerance** command. [See ecmp-alb.]

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Example: Configuring Aggregated Ethernet Load Balancing.]

- **VLAN TCC encapsulation on aggregated Ethernet interfaces (MX Series)**—Starting in Junos OS Release 20.1R1, aggregated Ethernet interfaces support VLAN translational cross-connect (TCC) encapsulation. For configuring VLAN TCC encapsulation, you must have the member links of aggregated Ethernet with VLAN TCC encapsulation supported hardware.

  > **NOTE:** MX Series routers do not perform any external commit check for member links of aggregated interfaces for the VLAN TCC encapsulation supported hardware.

  - Enable the **extended-vlan-tcc** option for aggregated Ethernet interfaces at the **[edit interfaces *interface-name* encapsulation]** hierarchy level to configure extended 802.1q tagging for TCC.
  - Enable the **vlan-tcc** option for aggregated Ethernet interfaces at the **[edit interfaces *interface-name* unit logical-unit-number encapsulation]** hierarchy level to configure 802.1q tagging for TCC.
  - Enable the **inet-address** option for aggregated Ethernet interfaces at the **[edit interfaces *interface-name* unit logical-unit-number family tcc proxy]** hierarchy level to configure proxy host address on the non-Ethernet side of Ethernet TCC circuits.
  - Enable the **inet-address** option for aggregated Ethernet interfaces at the **[edit interfaces *interface-name* unit logical-unit-number family tcc remote]** hierarchy level to configure remote host address on the non-Ethernet side of Ethernet TCC circuits.
  - Enable the **mac-address** option for aggregated Ethernet interfaces at the **[edit interfaces *interface-name* unit logical-unit-number family tcc remote]** hierarchy level to configure remote MAC address on the non-Ethernet side of Ethernet TCC circuits.
  - Enable the **tcc** option for aggregated Ethernet interfaces at the **[edit interfaces *interface-name* unit logical-unit-number family]** hierarchy level to configure the TCC protocol suite.

  [See Configuring VLAN TCC Encapsulation.]

- **MX2K-MPC11E supports Junos node slicing (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, the MX2K-MPC11E supports Junos node slicing and abstracted fabric (af) interfaces. Using

Junos node slicing, you can create multiple partitions in a single physical MX Series router. Each partition, referred to as a guest network function (GNF), behaves as an independent router. An af interface is a pseudointerface that exhibits a first-class Ethernet interface behavior. The af interface facilitates routing control and management traffic between GNFs through the switch fabric. In a Junos node slicing deployment, the MX2K-MPC11E interoperates with all MPCs that support the af interfaces.

> **NOTE:**
> - The MX2K-MPC11E interoperates only with the Switch Fabric Board SFB3.
> - The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Understanding Junos Node Slicing.]

- **Support for rate selectability on MX2K-MPC11E line cards (MX2010 and MX2020)**—In Junos OS Releases 19.3R2 and 20.1R1, we introduce a new fixed-configuration, rate-selectable line card, MX2K-MPC11E. The line-card has 40 built-in ports that can operate at 100-Gbps speed. You can configure all ports in a PIC to operate at the same speed or configure all the ports at different supported speeds. With QSFP28 optics installed, all ports operate at a default speed of 100-Gbps. In addition, you can use QSFP+ optics on Port 0 of every PIC and configure it as:
  - A 40-Gbps interface
  - Four 10-Gbps interfaces (channels), using breakout cables

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Introduction to Rate Selectability.]

- **Distributed LACP support in PPM AFT on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, the MX2K-MPC11E line card supports distributed LACP. Distributed LACP support is managed by the advanced forwarding toolkit (AFT)-based periodic packet manager (PPMAN). In earlier releases, and for other line cards except MPC10E, distributed LACP support is managed by the Junos OS-based PPMAN.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Periodic Packet Management.]

- **Optimize fabric path to prevent traffic hop with MX2K-MPC11E line cards (MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 20.1R1, on MX2008, MX2010, and MX2020 routers with MX2K-MPC11E, you can optimize the fabric path of the traffic flowing over abstracted fabric (af) interfaces between two guest network functions (GNFs) by configuring fabric optimization mode. This feature reduces fabric bandwidth consumption by preventing any additional fabric hop (switching of traffic flows from one Packet Forwarding Engine to another because of load balancing on the af interface) before the packets eventually reach the destination Packet Forwarding Engine.

  To configure fabric optimization mode, use the following CLI command at the base system (BSYS): **set chassis network-slices guest-network-functions gnf *id* collapsed-forward (monitor | optimize)**.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Optimizing Fabric Path for Abstracted Fabric Interface.]

- **Chassis and power management for MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, the MX2010 and MX2020 routers with the MX2K-MPC11E line card support chassis management features, including field- replaceable unit (FRU) management, power budgeting and management, and environmental monitoring.

  The MX2K-MPC11E line card supports the following configuration:

  - The ambient temperature is less than 46°C.
  - The ports on the MX2K-MPC11E line cards operate at various modes or speeds (10-Gbps, 40-Gbps, or 100 Gbps). The pic-mode specifies the speed of the active ports. If pic-mode is not specified, then the default mode is 100 Gbps.
  - Supports dynamic power management.
  - Supports both hyper mode (the default mode) and normal mode.
  - Supports both normal mode (the deafult mode) and enhanced priority mode for interface schedulers.
  - Supports interface queueing modes, namely WAN port queueing mode (the default mode), limited queueing mode, and enhanced queueing mode.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Understanding How Configuring Ambient Temperature Helps Optimize Power Utilization and Understanding How Dynamic Power Management Enables Better Utilization of Power.]

- **MPC Protocol and Application Support for MX2K-MPC11E line cards**—Starting in 20.1R1, MX2020 and MX2010 routers with MX2K-MPC11E line cards support many MPC protocols and applications. For a complete list, see *Protocols and Applications Supported by the MX2K-MPC11E*.
  - Standard Generic Routing Encapsulation (GRE)
  - Bidirectional Forwarding Detection protocol (BFD)
  - Internet Control Message Protocol (ICMP) and ICMPv6
  - Border Gateway Protocol (BGP)
  - BGP/MPLS virtual private networks (VPNs)
  - Logical system and Virtual routing and forwarding (VRF) routing instances
  - Load Balancing
  - Class of Service (CoS)—per port, virtual LAN (VLAN), Point-to-Point Protocol over Ethernet (PPPoE) or Dynamic Host Configuration Protocol (DHCP), Egress hierarchical class-of-service (CoS) shaping
  - Layer 2 Features
  - Firewall filters and policers

    NOTE: The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See MX Series 5G Universal Routing Platform Interface Module Reference.]

- **Support for new show | display set CLI commands (ACX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the following new **show** commands have been introduced:
  - **show | display set explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the top level of the hierarchy.
  - **show | display set relative explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the current hierarchy level.

  [See show | display set and show | display set relative.]

Junos OS, XML, API, and Scripting

- **The jcs:load-configuration template supports loading the rescue configuration (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **jcs:load-configuration** template supports the **rescue** parameter to load and commit the rescue configuration on a device. SLAX and XSLT scripts can call the **jcs:load-configuration** template with the **rescue** parameter set to **"rescue"** to replace the active configuration with the rescue configuration.

  [See Changing the Configuration Using SLAX and XSLT Scripts and jcs:load-configuration Template.]

Junos Telemetry Interface

- **IS-IS adjacency and LSDB event streaming support on JTI (MX960, PTX1000, and PTX10000)**—Junos OS Release 20.1R1 provides IS-IS adjacency and link-state database (LSDB) statistics using Junos telemetry interface (JTI) and remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. ON_CHANGE statistics are sent to an outside collector.

  The following resource paths are supported:

  - **/network-instances/network-instance/protocols/protocol/isis/interfaces/interfaces/circuit-counters/state/** (stream)
  - **/network-instances/network-instance/protocols/protocol/isis/interfaces/interface/levels/level/packet-counters/** (stream)
  - **/network-instances/network-instance/protocols/protocol/isis/levels/level/system-level-counters/state/** (stream)
  - **/network-instances/network-instance/protocols/protocol/isis/interfaces/interfaces/levels/level/adjacencies/adjacency/state/** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-ipv4-reachability/prefixes/prefix/state/** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-ipv4-reachability/prefixes/prefix/subtlvs/subtlv/** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv6-reachability/prefixes/prefix/state/** (ON_CHANGE)
  - **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv6-reachability/prefixes/prefix/subtlvs/subtlv/** (stream)
  - **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-is-reachability/neighbors/neighbors/subTLVs/subTLVs/adjacency-sid/sid/state/** (ON-CHANGE)
  - **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-is-reachability/neighbors/neighbors/subTLVs/subTLVs/lan-adjacency-sid/** (ON_CHANGE)

- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv4-interfaces-addresses/state/** (ON_CHANGE))
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv4-srlg/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv4-te-router-id/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv6-interfaces-addresses/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/router-capabilities/router-capability/subtlvs/subtlv/segment-routing-capability/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/state** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/area-address/state/address** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/nlpid/state/nlpid** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/lsp-buffer-size/state/size** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/hostname/state/hname** (stream)

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

- **Packet Forwarding Engine support for JTI on MX2K-MPC11E line cards (MX2010 and MX2020)**—Now supported in Junos OS Release 20.1R1, Junos telemetry interface (JTI) supports streaming of Packet Forwarding Engine statistics for MX2010 and MX2020 routers using Remote Procedure Calls (gRPC). gRPC is a protocol for configuration and retrieval of state information. This support was first introduced in Junos OS Release 19.3R2.

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

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

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

- **Platform, interface, and alarm sensor ON_CHANGE support on JTI (MX960, MX2020, PTX1000, PTX5000)**—Junos OS Release 20.1R1 supports platform, interface, and alarm statistics using Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services. You can use this feature to send ON_CHANGE statistics for a device to an outside collector.

  This feature supports the OpenConfig models:

  - **openconfig-platform.yang**: oc-ext:openconfig-version 0.12.1
  - **openconfig-interfaces.yang**: oc-ext:openconfig-version 2.4.1
  - **openconfig-alarms.yang**: oc-ext:openconfig-version 0.3.1

  Use the following resource paths in a gNMI subscription:

  - **/components/component** (for each installed FRU)
  - **/interfaces/interface/state/**
  - **/interfaces/interface/subinterfaces/subinterface/state/**
  - **/alarms/alarm/**

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

- **gRPC Dial-Out support on JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Junos OS Release 20.1R1 provides remote procedure call (gRPC) dial-out support for telemetry. In this method, the target device (server) initiates a gRPC session with the collector (client) and, when the session is established, streams the telemetry data that is specified by the sensor-group subscription to the collector. This is in contrast to the gRPC network management interface (gNMI) dial-in method, in which the collector initiates a connection to the target device.

  gRPC dial-out provides several benefits as compared to gRPC dial-in, including simplifying access to the target advice and reducing the exposure of target devices to threats outside of their topology.

  To enable export of statistics, include the **export-profile** and **sensor** statements at the [**edit services analytics**] hierarchy level. The export profile must include the reporting rate, the transport service (for example, gRPC), and the format (for example, gbp-gnmi). The sensor configuration should include the name of the collector (the server's name), the name of the export profile, and the resource path. An example of a resource path is **/interfaces/interface[name='fxp0'**.

  [See Using gRPC Dial-Out for Secure Telemetry Collection.]

- **gRPC version v1.18.0 with JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Junos OS Release 20.1R1 includes support for remote procedure call (gRPC) services version v1.18.0 with Junos telemetry interface (JTI). This version includes important enhancements for gRPC. In earlier releases, JTI is supported with gRPC version v1.3.0.

  Use gRPC in combination with JTI to stream statistics at configurable intervals from a device to an outside collector.

  [See gRPC Services for Junos Telemetry Interface.]

- **SR-TE statistics for uncolored SR-TE policies streaming on JTI (MX Series, PTX Series)**—Junos OS Release 20.1R1 provides segment routing traffic engineering (SR-TE) per label-switched Path (LSP) route statistics using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. Using JTI and gRPC services, you can stream SR-TE telemetry statistics for uncolored SR-TE policies to an outside collector.

  Ingress statistics include statistics for all traffic steered by means of an SR-TE LSP. Transit statistics include statistics for traffic to the Binding-SID (BSID) of the SR-TE policy.

  To enable these statistics, include the **per-source per-segment-list** statement at the [**edit protocols source-packet-routing telemetry statistics**] hierarchy level.

  If you issue the **set protocols source-packet-routing telemetry statistics no-ingress** command, ingress sensors are not created.

  If you issue the **set protocols source-packet-routing telemetry statistics no-transit** command, transit sensors are not created. Otherwise, if BSID is configured for a tunnel, transit statistics are created.

  The following resource paths (sensors) are supported:

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

  To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC.

  Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

  [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface, source-packet-routing, and show spring-traffic-engineering lsp detail name *name*.]

- **LLDP statistics, notifications, and configuration model for suppress-tlv-advertisement support on JTI (MX240, MX480, MX960, MX10003, PTX10008, PTX10016)**—Junos OS Release 20.1R1 provides remote procedure call (gRPC) streaming services support for attribute leaf **suppress-tlv-advertisement** under the resource path **/lldp/state/suppress-tlv-advertisement**. The following TLVs are supported, which in turn support operational state, notifications, and configuration change support:

  - port-description
  - system-name
  - system-description
  - system-capabilities
  - management-address
  - port-id-type

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

- **CPU and NPU sensors support using JTI on MX2K-MPC11E line cards (MX2010 and MX2020)**—Junos OS Release 20.1R1 supports Junos telemetry interface (JTI) CPU and network processing unit (NPU)

sensors on MX Series routers with MX2K-MPC11E line cards. JTI enables streaming statistics from these sensors to outside collectors at configurable intervals using remote procedure call (gRPC) services.

Unlike the Junos kernel implementation in earlier Junos OS releases that support these sensors, this feature uses the OpenConfig AFT model. Because of this, there is a difference in the resource path and key-value (kv) pair output compared to the Junos kernel output.

Use the following resource paths to export statistics:

/junos/system/linecard/cpu/memory/

/junos/system/linecard/npu/memory/

/junos/system/linecard/npu/utilization/

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

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

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

- **gNMI extension compliance with JTI (MX Series)**—Starting in Junos OS Release 20.1R1, changes are qualified in the extension header for Junos telemetry interface (JTI), ensuring they are compliant with the OpenConfig gnmi.extensions.proto specification.

  See gnmi-extensions.md.]

- **gNMI-based streaming telemetry support for Packet Forwarding Engine sensors on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, gRPC Network Management Interface (gNMI) service support is available to export Packet Forwarding Engine statistics for telemetry monitoring and management using Junos telemetry interface (JTI). Using gNMI and JTI, data is exported from devices to outside collectors at configurable intervals. This feature includes support (SensorD daemon) to export telemetry data for the OpenConfig model called AFT platform.

  Use the following resource paths to export sensor data for interface information and traffic, logical interface traffic, firewall filter counters, and policer counters:

  - **/junos/system/linecard/interface/**
  - **/junos/system/linecard/interface/traffic/**
  - **/junos/system/linecard/interface/queue/**
  - **/junos/system/linecard/interface/logical/usage/**

- **/junos/system/linecard/firewall/**
- **/junos/system/linecard/services/inline-jflow/**

To provision the sensor to export data through gNMI services, use the **Subscribe** RPC. The **Subscribe** RPC and subscription parameters are defined in the gnmi.proto file. Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

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

Layer 2 Features

- **Supported Layer 2 features on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, the following advanced Layer 2 features are supported on MX2K-MPC11E line cards:
  - **Forwarding CoS (Q-depth monitoring)**—You can configure a Junos telemetry interface sensor that exports queue depth statistics for egress queue traffic. Telemetry data is exported directly from the line card. You can also apply one or more regular expressions to filter data. Only UDP streaming of data is supported. gRPC streaming of queue depth statistics is not currently supported. [See sensor (Junos Telemetry Interface).]
  - **Layer 2 firewall forwarding support**. [See Layer 2 Port Mirroring Firewall Filters.]
  - **Layer 2 forwarding**—IRB, VLAN handling, and Q-in-Q tunneling. [See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation.] Virtual private LAN services (VPLS). [See Introduction to VPLS.] Firewall filters for Layer 2 and MAC filters. [See Layer 2 Forwarding Tables.]
  - **Multicast features**—P2MP (RSVP-TE P2MP and multipoint LDP inband) and P2MP interface support for PIM, Rosen multicast VPNs, and multicast-only fast reroute (MoFRR). [See Multicast Overview.]

    > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

- **Support for Layer 2 services with PWHT on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, some of the Layer 2 services are supported with pseudowire headend termination (PWHT) on the new MX2K-MPC11E line card.

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Protocols and Applications Supported by the MX2K-MPC11E and Layer 2 VPNs and VPLS User Guide for Routing Devices.]

- **Support for basic Layer 2 features on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, MX2010 and MX2020 routers with the MX2K-MPC11E line card supports the following basic Layer 2 features:
  - Layer 2 bridging with trunk and access modes
  - MAC learning and aging
  - Handling BUM (broadcast, unknown unicast and multicast) traffic, including split horizon
  - MAC move
  - Layer 2 forwarding and flooding statics
  - Mesh groups
  - Static MAC addresses
  - MAC learning and forwarding on AE interfaces
  - Bridging on untagged interfaces
  - Basic Q-n-Q tunneling (without VLAN-translation and VLAN map operations)

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Understanding Layer 2 Bridge Domains, Understanding Layer 2 Learning and Forwarding.]

Layer 3 Features

- **Support for new MPC11E line card (MX Series)**—Starting in Junos OS Release 20.1R1, we've introduced a new MPC, MPC11E, that supports the following Layer 3 features:

  The following Layer 3 features are supported on MPC11E in 20.1R1:

  - BGP
  - IS-IS
  - Layer 3 VPN
  - OAM - LSP/VPN ping, traceroute, automatic bandwidth, and MPLS-FRR link node protection

- OSPF
- RIP
- Tunnel (GRE tunnels, logical tunnels, and virtual tunnels)

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

- **Support for IPv6 Ping, IPv6 Traceroute and ECMP traceroute for Labelled-ISIS Segment Routing paths (MX Series and VMX)**— Starting in Release 20.1R1, Junos OS supports IPv6 Ping, IPv6 Traceroute, and equal-cost multipath (ECMP) traceroute for Labelled-ISIS Segment Routing paths.

Management

- **Error recovery, fault handling, and resiliency support for MX2K-MPC11E (MX2010 and MX2020)**—In Junos OS Releases 19.3R2 and 20.1R1, the MX2010 and MX2020 routers with MX2K-MPC11E line cards support error recovery, fault handling, and software resiliency. The MX2K-MPC11E line cards support detecting errors, reporting them through alarms, and triggering resultant actions. To view application level errors, use the **show trace node fpc<#> application fabspoked-pfe-redbull** command. To check the status of the card, use the **show chassis fpc pic-status** command. Use the **show chassis errors active** command to view the fault details and the **show system alarm** command to view the alarm details.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See show chassis fpc pic-status and clear chassis fpc errors.]

MPLS

- **Support for MPLS features on MX2K-MPC11E line card (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, the new MX2K-MPC11E line card supports some of the MPLS features.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Protocols and Applications Supported by the MX2K-MPC11E.]

- **Support for selective MPLS traffic mirroring (MX Series with MPC10)**—Starting in Junos OS Release 20.1R1, MX Series routers with MPC10 line cards supportelective MPLS traffic mirroring. You can apply inbound and outbound filters for the MPLS family based on MPLS-tagged IPv4 and IPv6 parameters

using inner payload match conditions, and enable selective port mirroring of MPLS traffic on to a monitoring device.

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

- **Support for segment routing over RSVP forwarding adjacency (MX Series and PTX Series)**—Starting with Junos OS Release 20.1R1, we provide support for segment routing traffic to be carried over RSVP LSPs that are advertised as forwarding adjacencies in IS-IS. This feature is implemented in a network having LDP on the edge and RSVP in the core where you can easily replace LDP with IS-IS segment routing because it eliminates the need for MPLS signaling protocols such as LDP. This helps to remove a protocol from the network and results in network simplification.

  [See Understanding Segment Routing over RSVP Forwarding Adjacency in IS-IS.]

- **Support for static adjacency segment identifier for aggregated Ethernet member links on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting with Junos OS Release 20.1R1, you can configure a transit single-hop static label-switched path (LSP) for a specific member link of an aggregated Ethernet (ae) interface. The label for this route comes from the segment routing local block (SRLB) pool of the configured static label range. Configure the aggregated Ethernet member-interface name using the **member-interface** statement option at the **[edit protocols mpls static-label-switched-path *name* transit *name*]** hierarchy level. This feature is supported for aggregated Ethernet interfaces only.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See transit and Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-Hop Static LSP.]

- **Support for Seamless MPLS Layer 3 features on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, MX2010 and MX2020 routers with the MX2K-MPC11E line card support the following MPLS Layer 3 features:

  - Redundant logical tunnel interfaces.
  - Pseudowire subscriber interfaces using either logical tunnel or redundant logical tunnel interfaces as anchor point.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Redundant Logical Tunnels Overview and MPLS Pseudowire Subscriber Logical Interfaces.]

- **Support for segment routing (SR) and segment routing traffic engineering (SRTE) statistics on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, MX2010 and

MX2020 routers with the MX2K-MPC11E line card supports segment routing (SR) and segment routing traffic engineering (SRTE) statistics.

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Understanding Source Packet Routing in Networking (SPRING).]

- **CoS-based forwarding and policy-based routing to steer selective traffic over an SR-TE path (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 20.1R1, you can use CoS-based forwarding (CBF) and policy-based routing (PBR, also known as filter-based forwarding or FBF) to steer service traffic using a particular segment routing-traffic-engineered (SR-TE) path. This feature is supported only on non-colored segment routing LSPs that have the next hop configured as a first hop label or an IP address.

  With CBF and PBR, you can:

  - Choose an SR-TE path on the basis of service.
  - Choose the supporting services to resolve over the selected SR-TE path.

  [See Example: Configuring CoS-Based Forwarding and Policy-Based Routing For SR-TE LSPs.]

- **Support for MPLS features on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, some of the MPLS features are supported on the new MX2K-MPC11E line card.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Protocols and Applications Supported by the MX2K-MPC11E.]

Multicast

- **Support for multicast forwarding on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, multicast forwarding is fully supported on MX2010 and MX2020 routers with MX2K-MPC11E line cards and Enhanced Switch Fabric Boards (SFB3).

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Multicast Overview.]

- **Next-generation multicast VPN supported on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, the MX2K-MPC11E line card supports next-generation MVPN.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Multicast Overview.]

Network Management and Monitoring

- **On-box monitoring support on the control plane (MX Series and PTX Series)**—Starting in Junos OS Release 20.1R1, you can configure on-box monitoring to monitor anomalies with respect to the memory utilization of Junos OS applications and the overall system in the control plane of MX Series and PTX Series routers.

  You can use on-box monitoring to monitor system-level memory and process-level memory to detect possible leaks. When the system is running low on memory, the process heuristic shares the prediction and you can configure the action to be taken when leaks are identified.

  See memory (system)

- **Enhanced PKI traps, log notifications, and SNMP for IPsec VPN (MX Series with USF and the SRX5000 line of devices with SPC3 card)**—Starting in Junos OS Release 20.1R1, you can enable the peer down and IPsec tunnel down traps and configure the certificate authority (CA) and local certificate traps. We've enhanced the existing IPsec VPN flow monitor MIB **jnxIpSecFlowMonMIB** to support the global data plane, active IKE SA, active IPsec SA, and active peer statistics for tunnels using IKEv2. We've also enhanced the output of the **show security ike stats** command to add additional options **(<brief> | <detail>)**. Use the **clear security ike stats** command to clear the IKEv2 statistic counters.

  [See Configure the Certificate Expiration Trap, Enterprise-Specific SNMP MIBs Supported by Junos OS, Enable Peer Down and IPsec Tunnel Down Traps, trap (Security PKI), trap (Security IKE), clear security ike stats, show security ike stats, show security ipsec statistics, show security ike security-associations, and show security ike active-peer.]

Next Gen Services

- **Support for Port Control Protocol (PCP)**—Starting in Junos OS Release 20.1R1, Next Gen Services supports the Port Control Protocol (PCP), which provides a mechanism to control how incoming packets are forwarded by upstream devices such as Network Address Translator IPv6/IPv4 (NAT64), Network Address Translator IPv4/IPv4 (NAT44), and IPv6 and IPv4 firewall devices, and mechanism to reduce application keep alive traffic.

  [See pcp-rules.]

- **Support for Traffic Load Balancer**—Starting in Junos OS Release 20.1R1, Next Gen Services support Traffic Load Balancer (TLB). TLB enables you to distribute traffic among multiple servers.

[See Traffic Load Balancer Overview.]

- **Support for TLS transport protocol for Next Gen Services CGNAT syslog messages**—Starting in Junos OS Release 20.1R1, you can configure the transport security protocol for Next Gen Services CGNAT global syslog messages to Transport Layer Security (TLS), as well as UDP or TCP.

  TLS ensures the secure transmission of data between a client and a server through a combination of privacy, authentication, confidentiality, and data integrity. SSL relies on certificates and private-public key exchange pairs for this level of security.

  [See transport.]

- **Next Gen Services on GNFs (MX480 and MX960)**—Starting in Junos OS Release 20.1R1, guest network functions (GNFs) on MX480 and MX960 routers support Next Gen Services when the MX-SPC3 Services Processing Card is installed. You can enable Next Gen Services on a GNF by using the existing command **request system enable unified-services** at the GNF level. In a Junos node slicing setup, you can use both MX-SPC3 and MS-MPC on the same chassis but on different GNFs. However, the MX-SPC3 comes online only if you have enabled Next Gen Services on the GNF. If you have not enabled Next Gen Services, only the MS-MPC comes online.

  > **NOTE:** The MX-SPC3 does not support abstracted fabric interfaces.

  [See Enabling and Disabling Next Gen Services and request system enable unified-services.]

- **Support for URL filtering, DNS sinkhole and Juniper Sky ATP URL filtering** —Starting in Junos OS Release 20.1R1, under Next Gen Services you can configure DNS filtering to identify DNS requests for blacklisted website domains and URL filtering to determine which Web content is not accessible to users. We also support Juniper Sky ATP filtering, which is a cloud-based solution that integrates with Policy Enforcer on the Junos Space Security Director.

  [See local-category.]

OAM

- **Support for link fault management (MX2K-MPC11E)**—Starting in Junos OS Release 20.1R1, you can configure IEEE 802.3ah link fault management (LFM) for MX2K-MPC11E on MX2010 and MX2020 routers. You can also configure the following supported LFM features:

  - Discovery and link monitoring
  - Distributed LFM
  - Remote fault detection and remote loopback

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Introduction to OAM Link Fault Management (LFM).]

Port Security

- **Media Access Control Security (MACsec) support on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, MACsec is supported on MX2010 and MX2020 routers with the MX2K-MPC11E line card. MACsec is an industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. The MPC11E supports MACsec on all 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet interfaces. The supported cipher suites are GCM-AES-256 and GCM-AES-128. Only static CAK mode is supported.

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

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

- **VLAN-level MACsec with unencrypted VLAN tags (MX10003 with JNP-MIC1-MACSEC)**—You can establish MACsec sessions for logical interfaces instead of physical interfaces on MX10003 routers with the JNP-MIC1-MACSEC installed. VLANs tags are now transmitted in cleartext, allowing intermediate switches that are MACsec-unaware to process VLAN tags. This feature enables MACsec encryption of point-to-multipoint VLAN connections over service provider WANs.

  [See Media Access Control Security (MACsec) over WAN.]

Routing Policy and Firewall Filters

- **Support for CCC and Layer 3 firewall forwarding on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting with Junos OS Release 20.1R1, circuit cross-connect (CCC) traffic and Layer 3 firewall forwarding features are supported on MX2K-MPC11E line cards.

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See CCC Overview.]

- **Support for firewall forwarding on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, firewall forwarding is fully supported on MX2010 and MX2020 routers with MX2K-MPC11E line cards and Enhanced Switch Fabric Boards (SFB3s).

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Filter-Based Forwarding Overview.]

- **Support for IPv6 discard interfaces (MX Series)**—Starting in Junos OS Release 20.1R1, you can configure a discard interface for IPv6 traffic. Do this at the **[edit interfaces dsc unit 0 family inet6]** hierarchy level.

  [See Configuring Discard Interfaces]

Routing Protocols

- **Support for topology-independent loop-free alternate (TI-LFA) in IS-IS for IPv6-only networks (ACX Series, MX Series, and PTX Series)**— Starting with Junos OS Release 20.1R1, you can configure TI-LFA with segment routing in an IPv6-only network for the IS-IS protocol. TI-LFA provides MPLS fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. TI-LFA provides protection against link failure and node failure.

  You can enable TI-LFA for IS-IS by configuring the **use-post-convergence-lfa** statement at the **[edit protocols isis backup-spf-options]** hierarchy level. You can enable the creation of post-convergence backup paths for a given IPv6 interface by configuring the **post-convergence-lfa** statement at the **[edit protocols isis interface *interface-name* level *level*]** hierarchy level. The **post-convergence-lfa** statement enables link-protection mode.

  You can enable **node-protection** mode for a given interface at the **[edit protocols isis interface *interface-name* level *level* post-convergence-lfa]** hierarchy level. However, you cannot configure fate-sharing protection for IPv6-only networks.

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

- **Support for IP forward backup path for BGP-LS peer SIDs (MX Series)**— Starting in Junos OS Release 20.1R1, 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 LS updates. If you have configured both segment protection and IP forwarding backup path, then backup segment protection takes precedence over the IP forwarding 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.]

Services Applications

- **Support for port mirroring (MX2K-MPC11E line card on MX2010 and MX2020 routers)**—In Junos OS Releases 19.3R2 and 20.1R1, you can configure port mirroring on the MX2K-MPC11E line card to mirror a copy of a packet to a configured destination, in addition to the normal processing and forwarding of the packet. The MX2K-MPC11E supports IPv4 (inet) and IPv6 (inet6) address families only.

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Configuring Port Mirroring.]

- **Support for tunnel interfaces (MX2K-MPC11E line card on MX2010 and MX2020 routers)**—In Junos OS Releases 19.3R2 and 20.1R1, Junos OS supports three tunnel interfaces: generic routing encapsulation (GRE) tunnel, logical tunnel (LT), and virtual tunnel (VT) on the MX2K-MPC11E line card.

  - 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 **family inet**, **family inet6**, and **family iso** options. The **encapsulation** statement supports the Ethernet and VLAN physical interface options only.
  - The VT interface supports the **family inet** option only.

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Tunnel Services Overview.]

- **Fabric support on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, the MX2K-MPC11E line card is introduced. It is composed of 8 Packet Forwarding Engines per FPC. Each Packet Forwarding Engine on the MX2K-MPC11E line card has 3 fabric planes per SFB, which is a total of 24 fabric planes. All Packet Forwarding Engines have fabric connectivity with the SFB3. The fabric links are monitored for cyclic redundancy check (CRC) errors. Each Packet Forwarding Engine supports 500G fabric throughput when all 24 fabric planes are operational.

  NOTE:
  - Fabric redundancy is not supported on MX2K-MPC11E line card. The MX2K-MPC11E line card interoperates only with SFB3.
  - The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Protocols and Applications Supported by the MX2K-MPC11E.]

- **Support for local preference when selecting forwarding next hops for load balancing on MPC11E (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, you can have traffic flows across aggregated Ethernet or redundant logical-tunnel interfaces prefer local forwarding next hops over remote ones, for example to ensure that the overall load on the fabric is reduced.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See local-bias.]

- **Inline J-Flow support for EVPN traffic (MX Series with MPC10 and MPC11)**—Starting with Junos OS Release 20.1R1, you can use inline J-Flow sampling for the bridge family. You can monitor Inline J-Flow traffic hitting the bridge family and report the necessary fields in either Version 9 or IPFIX format. The new bridge family under the **forwarding-options sampling instance** hierarchy monitors all traffic hitting the VPLS or bridge family.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Understanding Inline Active Flow Monitoring.]

- **Configure next-hop-based dynamic tunnels on MX2K-MPC11E line card (MX2010 and MX2020 routers)**—In Junos OS Releases 19.3R2 and 20.1R1, on MX2010 and MX2020 routers with an MX2K-MPC11E line card, you can configure next-hop-based dynamic tunnels for the following configurations:

  - MPLS-over-UDP—You can configure a dynamic MPLS-over-UDP tunnel that includes a tunnel composite next hop.

    In a dynamic tunnel configuration, where the Routing Engine forwards a lot of routes to the Packet Forwarding Engine, the FIB convergence may take more time resulting in traffic loss. Also, when you restart an FPC restart in a dynamic tunnel configuration, traffic flow may not resume.

  - MPLS-over-GRE—You can configure MPLS LSPs to use generic routing encapsulation (GRE) tunnels to cross routing areas, autonomous systems, and ISPs.

    > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See dynamic-tunnels.]

- **Support for inline active flow monitoring on (MPC11E line cards on MX240, MX480, and MX960)**—Starting in Junos OS Release 20.1R1, you can perform inline flow monitoring to support:
  - MPLS, MPLS-IPv4, and MPLS-IPv6
  - IPv4 or IPv6 traffic on next-hop based GRE tunnels and ps interfaces

  Both IPFIX and V9 formats are supported.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Understanding Inline Active Flow Monitoring.]

- **Support for Two-Way Active Measurement Protocol (TWAMP) on MX2K-MPC11E line card (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.1R1, the MX2K-MPC11E line card supports TWAMP. You can use the TWAMP-Control protocol to set up performance measurement sessions between a TWAMP client and a TWAMP server, and use the TWAMP-Test protocol to send and receive performance measurement probes. Configuring the TWAMP client instance to use si-x/y/z as the destination interface (which enables inline services) is not supported if the router has an MX2K-MPC11E installed in the chassis. You can configure only the none authentication mode on the line card.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

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

- **L2TPv2 silent failover on peer interface for L2TPv2 subscriber services on MX2K-MPC11E line card (MX2010 and MX2020 routers)**—In Junos OS Releases 19.3R2 and 20.1R1, you can configure L2TPv2 silent failover and peer interface support for L2TPv2 subscriber services on MX2K-MPC11E line card.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Peer Resynchronization After an L2TP Failover.]

- **Port mirroring support on MX2K-MPC11E (MX2010 and MX2020 routers)**—In Junos OS Releases 19.3R2 and 20.1R1, you can configure port mirroring on the MX2K-MPC11E line card. You can configure port mirroring for the CCC, bridge, and family any only.

> **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Understanding Port Mirroring.]

- **FlowTapLite support on MPC10E (MX240, MX480, and MX960 routers)**—Starting in Junos OS Release 20.1R1, you can configure FlowTapLite on the MPC10E line card.

  [See Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs.]

- **Support adaptive load balancing (ALB) for ECMP next hops (MX Series)**—Currently, adaptive load balancing for ECMP next hops is limited to a single Packet Forwarding Engine. Hence, traffic is restricted to a single Packet Forwarding Engine and impacts the flexibility and redundancy. Starting in Junos OS Release 20.1R1, you can configure adaptive load balancing for ECMP next hops across multiple ingress Packet Forwarding Engines on the same line card for even distribution of the traffic and redundancy. The behavior is default starting with Junos OS Release 20.1R1 and you cannot choose to configure back to the behavior prior to Junos OS Release 20.1R1. Also, the behavior is not applicable when you configure adaptive load-balancing and locality-bias together.

  To configure adaptive load balancing for ECMP next hops, configure the **ecmp-alb** command under the **[edit chassis]** hierarchy level.

  [See ecmp-alb].

Software Defined Networking

- **Delegate segment routing LSPs to a PCE (MX Series)**—Starting in Junos OS Release 20.1R1, you can enable a Path Computation Client (PCC) to delegate locally configured IPv4 non-colored segment routing LSPs to a Path Computation Element (PCE) controller. The PCE controls the delegated LSPs and can modify LSP attributes for traffic steering.

  A PCC with delegation capability can take back control of the delegated segment routing LSPs from the PCE when the PCEP session goes down; the LSPs would otherwise be deleted from the PCC. You can thus ensure LSP data protection by averting a situation where packets are silently discarded or dropped (also known as a traffic black-hole condition).

  [See Segment Routing for the Path Computation Element Protocol Overview and Example: Configuring Path Computation Element Protocol for SPRING-TE LSPs.]

Subscriber Management and Services

- **Support for BNG M:N subscriber redundancy over pseudowire interfaces (MX Series)**—Starting in Junos OS Release 20.1R1, you can configure BNG M:N redundancy using pseudowire redundancy in addition to using VRRP redundancy. The pseudowire redundancy method is supported for IP/MPLS network and Layer 2 VPN scenarios using pseudowire tunnels. These scenarios support dynamic N:1 VLANs.

  [See M:N Subscriber Redundancy Overview.]

- **Distributed denial of service protection on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, the MX2K-MPC11E line cards support DDoS protection.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Protocols and Applications Supported by the MX2K-MPC11E.]

- **Subscriber services uplink support on MX2K-MPC11E line cards (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, you can use the MX2K-MPC11E line cards for uplink connections to the core network. This support requires you to enable enhanced subscriber management.

  > **NOTE:** The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

  [See Protocols and Applications Supported by the MX2K-MPC11E.]

- **Support for managing policy and charging rules function (PCRF) server errors (MX Series)**—Starting in Junos OS Release 20.1R1, you can configure the router to reinitialize the PCRF session when triggered by certain PCRF server errors that result in a state mismatch between the server and the router. You

can also configure the router to generate an extended session ID that is universally unique by appending a 32-bit session stamp based on the current UTC time when the router creates the CCR-GX-I.

[See Understanding Gx Interactions Between the Router and the PCRF.]

System Management

- **Precision Time Protocol (PTP) and IRB support on MPC7E line cards (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.1R1, we support PTP over IRB on master interface configurations for MPC7E line cards. This release also supports the configuration of aggregated Ethernet over IRB. We've also added **disable-lag-revertive-switchover** statement at a global level. This configuration enables nonrevertive switchover for a LAG.

  NOTE:
  - Two-step clock mode is not supported.
  - PTP aggregated Ethernet child link switchover is not hitless, in both negotiated and nonnegotiated cases, in scenarios with aggregated Ethernet, because the client goes through a resynchronization phase. When unicast negotiation is enabled, the PTP backup clock starts fresh with new negotiation messages using the secondary link whenever the current active link goes down.
  - Aggregated Ethernet with mixed-speed child links is not supported over IRB.

  [See Configuring Precision Time Protocol Over Integrated Routing and Bridging.]

- **Restrict option under NTP configuration is now visible (ACX Series, QFX Series, MX Series, PTX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **noquery** command under the **restrict** hierarchy is now available and can be configured with a mask address. The **noquery** command is used to restrict ntpq and ntpdc queries coming from hosts and subnets.

  [See Configuring NTP Access Restrictions for a Specific Address.]

User Interface and Configuration

- **Synchronous Ethernet support for MPC11E (MX2010 and MX2020)**—Starting in Junos OS Release 20.1R1, Synchronous Ethernet is supported on the MPC11E.

  NOTE: Junos OS dose not support synchronous Ethernet clock recovery from MIC and Precision Time Protocol (PTP).

  NOTE: The MX2K-MPC11E line card is also supported in Junos OS Release 19.3R2 and later Junos OS 19.3 releases. It is not supported in any Junos OS 19.4 releases.

[See Synchronous Ethernet Overview.]

SEE ALSO

What's Changed

IN THIS SECTION

Learn about what changed in this release for MX Series routers.

What's Changed in Release 20.1R3

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.

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.]

- **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).]

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).]

MPLS

- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the cspf-backoff-time statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

  [See cspf-backoff-time..]

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.]

Routing Protocols

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the lsdist.0 and lsdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, we added multiple secondary loopback addresses in the traffic engineering database to the lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router ID.

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.]

What's Changed in 20.1R2

Class of Service (CoS)

- We've corrected the output of the **show class-of-service interface | display xml** command. Output of the following sort: <container> <leaf-1> data </leaf-1><leaf-2>data </leaf-2> <leaf-3> data</leaf-3> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container> will now appear correctly as <container> <leaf-1> data </leaf-1><leaf-2>data </leaf-2> <leaf-3> data</leaf-3></container> <container> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container>.

General Routing

- **Displaying accurate aggregate drop statistics (MX Series)**—You can view the accurate aggregate drop statistics when a packet drop is seen on an aggregated Ethernet interface by using the **show interfaces extensive** command. In earlier releases, the **show interfaces extensive** command did not display accurate aggregate drop statistics. Only the individual aggregate child interface displayed accurate drop statistics.
- **MS-MPC and MS-MIC service package (MX240, MX480, MX960, MX2020, MX2010, and MX2008)**—PICs of the Multiservices MPC (MS-MPC) and Multiservices MIC (MS-MIC) do not support any service package other than extension-provider. These PICs always come up with the extension-provider service package, irrespective of the configuration. If you try to configure any other service package for these PICs by using the command **set chassis fpc slot-number pic pic-number adaptive-services service-package**, an error is logged. Use the **show chassis pic fpc-slot *slot* pic-slot *slot*** command to view the service package details of the PICs of the MS-MPC and MS-MIC.

  [See extension-provider.[
- **Round-trip time load throttling for pseudowire interfaces (MX Series)**—The Routing Engine supports round-trip time load throttling for pseudowire (ps) interfaces. In earlier releases, only Ethernet and aggregated Ethernet interfaces are supported.

  [See Resource Monitoring for Subscriber Management and Services.[
- **Command to view summary information for resource monitor (MX Series routers and EX9200 line of switches)**—You can use the **show system resource-monitor** command to view statistics about the use of memory resources for all line cards or for a specific line card in the device. The command also displays

information about the status of load throttling, which manages how much memory is used before the device acts to reduce consumption.

[See show system resource-monitor and Resource Monitoring for Subscriber Management and Services.]

- **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/interfacesid='*sid-value*'/** 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 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.]

High Availability (HA) and Resiliency

- **IPv6 address in the prefix TIEs displayed correctly**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.
- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

Infrastructure

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

  [See interface-transmit-statistics.]

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**— You can set the verbosity of the trace log to only show error messages using the error option at the **[edit system services extension-service traceoptions[** hierarchy level.

  [See traceoptions (Services).]

Junos OS XML API and Scripting

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

  [See Junos XML API Explorer - Operational Tags.]

Network Management and Monitoring

- **Support for clearing the event at MEP level (MX Series)**—In Junos OS 20.1R2 and later, you can define an action profile for connectivity fault management at the local MEP level or at the remote MEP level. You define an action profile to monitor events and thresholds and specify an action that the device performs when the configured event occurs. When you define the action profile at the local MEP level, you can clear the event for the configured action profile at the local MEP level by specifying only the local MEP numeric identifier. When you define the action profile at the remote MEP level, you can clear the event for the configured action profile at the remote MEP level by specifying the local MEP numeric identifier as well as the remote MEP numeric identifier.

  [See clear oam ethernet connectivity-fault-management event.]

Routing Protocols

- **Automatic installation of YANG-based CLI for RIFT protocol (MX Series, QFX Series, and vMX with 64-bit and x86-based servers)**—In Rift 1.2 Release, installation of the CLI for RIFT protocol occurs automatically along with the installation of the junos-rift package. In the pre-1.0 releases of the junos-rift package, the RIFT CLI had to be installed separately using **request system yang** command after installation of the junos-rift package.–

- **Advertising /32 secondary loopback addresses to Traffic Engineering Database (TED) as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—In Junos OS Release, multiple loopback addresses export into lsdist.0 and lsdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router-ids instead of prefixes. In earlier Junos OS releases, multiple secondary loopback addresses in TED were added into lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router-id.

Services Applications

- **New option for configuring delay in IPsec SA installation**—In Junos OS Release 20.1R2, you can configure the **natt-install-interval *seconds*** option under the **[edit services ipsec-vpn rule *rule-name* term *term-name* then dynamic]** hierarchy to specify the duration of delay in installing IPsec SA in a NAT-T scenario soon after the IPsec SA negotiation is complete. The default value is 0 seconds.

Subscriber Management and Services

- **Improved tunnel session limits display (MX Series)**—Starting in Junos OS Release 20.1R2, the **show services l2tp tunnel extensive** command displays the configured value for maximum tunnel sessions. On both the LAC and the LNS, this value is the minimum from the global chassis value, the tunnel profile value, and the value of the Juniper Networks VSA, Tunnel-Max-Sessions (26–33). On the LNS, the configured host profile value is also considered.

  In earlier releases, the command displayed the value 512,000 on the LAC and the configured host profile value on the LNS.

  [See Limiting the Number of L2TP Sessions Allowed by the LAC or LNS.]

What's Changed in Release 20.1R1

General Routing

- **Logical Interface is created along with physical Interface by default (EX Series switches, QFX Series switches, MX Series routers)**?The logical interface is created on ge, et, xe interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces were created. For example, for ge interfaces, earlier when you view the **show interfaces** command, by default, only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.
- **Precision Time Protocol (PTP) interface configuration (MX2020, MX2010, MX480, MX960, and MX240)**—Remove the aggregated Ethernet interface association and upgrade the device when configuring PTP interface.

Interfaces and Chassis

- **Displaying accurate aggregate drop statistics (MX Series)**—Starting in Junos OS Release 20.1R1, you can view the accurate aggregate drop statistics when a packet drop is seen on an aggregated Ethernet Interface by using the **show interfaces extensive** command. In earlier releases, the **show interfaces extensive** command did not display accurate aggregate drop statistics. Only the individual aggregate child interface displayed accurate drop statistics.

Network Management and Monitoring

- **Change in startup notification after GRES (MX Series routers)**– Starting in Junos OS Release 20.1R1, the master Routing Engine sends a **coldStart** notification when a device comes up. The master Routing Engine also sends **warmStart** notifications for subsequent restarts of the SNMP daemon. After graceful routing engine switchover (GRES) the new master Routing Engine sends a single **warmStart** notification and the backup Routing Engine does not send any notification. In earlier releases, after GRES, the new master RE would sometimes send two notifications or a single notification. Of these, the first notification was always a **coldStart** notification and the second was either a **coldStart** notification or a **warmStart** notification.
- **Enhancement to the show SNMP mib command**– In Junos OS Release 20.1R1, and later, a new option, **hex**, is supported to display the SNMP object values in the hexadecimal format. In earlier releases, the **show snmp mib** command displays the snmp object values in ASCII and decimal format only.

  See show snmp mib

Services Applications

- **Update to CLI option for configuring the version number to distinguish between currently supported version of the Internet draft draft-ietf-softwire-map-03**—In Junos OS Release 20.1R1, the **version-3** option under the **[edit services softwire softwire-concentrator map-e]** hierarchy for configuring the version number to distinguish between currently supported version of the Internet draft draft-ietf-softwire-map-03 is optional. In the earlier Junos OS releases, if you did not configure the **version-3** option, the configuration resulted in an error.

  [See map-e.]

Subscriber Management and Services

- **Single memory map applies to configuration and schema databases (MX Series)**—Starting in Junos OS Release 20.1R1, the Junos OS configuration database and the schema database share the same memory space. This means that when you set the maximum database size, the result is the total memory available to both of these databases. In earlier releases, the schema database is separate and fixed in size.

  [See Configuring Junos OS Enhanced Subscriber Management.]

SEE ALSO

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

- MX Series platforms with the FPC-PTX-P1-A or FPC2-PTX-P1A line card might encounter a single event upset (SEU) event that can cause a linked-list corruption of the TQ chip. The following syslog message is 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, a restart of the FPC is needed. Contact your Juniper support representative if the issue is seen after an FPC restart. PR1254415

- In some scenarios with MPC, the following major alarm and following messages are generated: **messages log: fpcx XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8) alarmd[3158]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC x Major Errors Major alarm set, FPC x Major Errors fpcx XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8) cli> show chassis alarms 1 alarms currently active Alarm time Class Description 2019-01-25 15:18:03 UTC Major FPC x Major.**

  Despite the major alarm set, this error is due to the Unknown Error Address logged in hardware to the DQ underrun. This message is harmless and has no service impact. PR1303489

- The MX104 router has the following limitations in error management:
  - The **show chassis fpc error** command is not available for MX104 in Junos OS Releases 13.3R7, 15.1R2,14.1R5,14.2R4, 13.3R8, and later.
  - Junos OS does not initiate restart of the system on encountering a fatal error.
  - Although you can configure **disable-PFE** for major errors action, Junos OS does not disable its only Packet Forwarding Engine on encountering a major error. PR1413314

- After an MX Series router with the JNP10K-LC2101 line card is powered on, a voltage of 1345–1348 mV is read for about 20 seconds, which gets stabilized to 1493 mV. During this period, the FPC x Voltage Tolerance Exceeded major alarm is raised. PR1415671

- The Routing Engine interprets any input from the console port as interrupts. Depending on the frequency, console noise impacts the Routing Engine interruption handling to different extents, even with the current mechanism. When the interrupt frequency is too high for the Routing Engine to handle, the impact might vary from the line card reboot (partial impact) to the Routing Engine reboot (chassis-wide impact). PR1436386

- On the MPC11E line card, the **number-of-sub-ports** configuration on the 4x10G channelized ports might cause the channels to go down. PR1442439

- In a scaled scenario where the Routing Engine pushes a lot of routes to the Packet Forwarding Engine in the presence of the dynamic tunnel configuration, FIB convergence might take more time, leading to traffic drops. PR1454817
- On the MPC11E line card, the following error message is seen when the line card is online: **i2c transaction error (0x00000002)**. PR1457655
- Dynamic SR-TE tunnels does not get automatically re-created at the new primary Routing Engine after the Routing Engine switchover. PR1474397
- Please enable Phytimestamping for PTP support. PR1480938
- If the frequency of messages is too high for the number of sessions, the kernel might be overloaded and causes riftd to quit and generate a core file. PR1481169
- After dot1xd restart, it takes around 60–70 seconds to display the MACsec statistics. In case of a full scale of 200 sessions, it takes around 40 seconds for all MACsec sessions to get configured on the FPC, then it reads statistics for 50 logical interfaces in every 5 seconds. This is to ensure that FPC CPU utilization is under check even on a scaled setup due to a slow MDIO bus. So, on a fully scaled setup of 200 logical interfaces, statistics for a given logical interface are read every 25 seconds. Thus the approximate total time will be 40 seconds + 25 seconds for statistics to be displayed after dot1x restart. PR1484699
- After an FPC restart or offline/online event, DUT sends the delete path for interfaces where LLDP is explicitly disabled. PR1484734
- The control peer PFCP heartbeat request timeout window must be greater than 90 seconds. PR1459135
- The traffic on GRE interface on both ingress and egress cannot be Layer 2 mirrored. PR1462375
- The following error message is issued when the connection between aftman and aft-ulcd is dropped: **[Error] aft-ipc: AFT-ULCD IPC: Program will exit - ERROR MESSAGE.** PR1467246
- If you move the MX2K-MPC11E line card from one GNF to another in an in-chassis Junos node slicing setup, the line card takes a longer time than expected to come online. PR1469729
- ALB over 64 links cannot rebalance the traffic to the desired configured tolerance. PR1470717
- When L2-Bridge domain is config with all four channelized interfaces into different Bridge Domain's of different VLAN IDs. Now sending traffic to one of the VLAN IDs - traffic flows correctly but statistics are sometime updated incorrectly on other three channelized interfaces. PR1472464
- On MX10003 and MX204 routers, BFD or LACP might flap during BGP convergence. PR1472587
- The following error message might appear: **Failed to complete DFE tuning**. This error message has no functional impact and can be ignored. PR1473280
- The aftd hogs on executing the clear VPLS table and MACs are not learned for less than 5 minutes. PR1473334
- ALB over 64 links fails to redistribute the traffic load after removing one fat traffic flow. PR1473435

Infrastructure

- The Juniper Routing Engine with HAGIWARA CF card installed, after upgrading to Junos OS Release 15.1 and later, the following error message might appear on the log: **smartd[xxxx]: Device: /dev/ada1, failed to read SMART Attribute Data** PR1333855

Interfaces and Chassis

- When Aggregate Ethernet interface with member links located on both MPC10/11 and non-MPC10/11 line cards, if it belongs to EVPN (Ethernet VPN) instance and configured CFM on it, due to difference in handling VLAN information between these two linecards, the CFM PDU (Protocol Data Unit) may be discarded. It will lead to the CFM session Down and the EVPN in Down state either. PR1543641

MPLS

- Process rpd might crash after network service configuration changed (like changing the range of MPLS labels) without rebooting all the Routing Engines (which is a system mandatory step). PR1461468
- On all platforms running Junos OS with distributed CSPF under SR-TE scenario, if you execute some operations such as deactivate or activate SR protocols, restart routing, and so on, then rpd crash might be observed. PR1493721
- With local reversion ON, there is a possibility of transit router not informing head-end of RSVP disabled link when link is flapped more than once. Work around is to remove local-reversion configuration. PR1576979

Platform and Infrastructure

- Traffic might drop due to the memory error of QX-chipset MPC. PR1197475
- Interface-group based firewall filters used at MX Series router with the VPLS and BRIDGE logical interfaces hosted by an MPC might work unpredictably. PR1216201
- Unknown unicast filter applied in an EVPN routing instance blocks unexpected traffic. PR1472511
- An EVPN does not support individual logical interfaces operation if ESI is configured on the physical interface. On MX Series routers, the loop prevention feature supports per logical interface flap, only when ESI is configured at the logical interface level. If one logical interface is flapped during BUM traffic flow through another logical interface, the loop prevention feature does not work. If ESI is configured on the physical interface and one of the ESI logical interfaces is down, EVPN considers it as whole physical interface down. The router's designated forwarder (DF) role or MH status is changed and the

local bias filter is reconfigured or deleted. Because the traffic flowing through the logical interface is not flapped, loop prevention is not enabled for the logical interface. PR1485100

- On all Junos OS platforms that support EVPN-MPLS or EVPN-VXLAN, when an existing ESI interface flaps or is newly added to the configuration, sometimes DF (Designated Forwarder) election happens before the local bias feature is enabled and during this time, existing broadcast, unknown unicast, and multicast (BUM) traffic might be looped for a short time duration (less than several seconds). PR1493650

VPNs

- In MVPN scenario with static lsp mapping type 3 route withdraw behavior can differ. PR1466122

Services Applications

- Currently, while configuring a DNS filter profile at the **[edit services web-filter profile *profile-name* dns-filter-template]** hierarchy level, you can configure a maximum of number of 32 DNS filter templates. However, for a profile configured under **[edit services web-filter profile *profile-name* security-intelligence-policy]** hierarchy level, you can configure more than 32 templates.

  [See dns-filter-template and security-intelligence-policy].

Subscriber Management and Services

- For dual-stacked clients over the same PPP-over-L2TP LNS session, enhanced subscriber management does not support configurations where both of the following are true:
  - The CPE sends separate DHCPv6 solicit messages for the IA_NA and the IA_PD.
  - The solicit messages specify a type 2 or type 3 DUID (link-layer address).

  As a workaround, you must configure the CPE to send a single solicit message for both IA_NA and IA_PD when the other configuration elements are present. PR1441801

SEE ALSO

What's New | **95**

What's Changed | **125**

Open Issues | **138**

Resolved Issues | **158**

Documentation Updates | **211**

Migration, Upgrade, and Downgrade Instructions | **212**

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

- The fxp0 marked as **Dest-route-down** because of specific operations such as disabling and enabling operations. PR1052725
- On a vMX instance, the performance of an X710 NIC is lower compared to the performance of an 82599 NIC. A 10-Gbps line rate can be achieved at a 512 byte packet size for the X710 NIC compared to 256 bytes for the 82599 NIC. PR1281366
- 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 on NFX150 devices to check the interface details, the system will not check whether the interface name provided is valid or invalid. The system will not generate an error message if the interface name is invalid. PR1306191
- The chain-composite statement does not bring in a lot of gain because TCNH is based on an ingress rewrite premise.PR1318984
- In a Message Queuing Telemetry Transport (MQTT) scenario, about 4000 KB of memory leakage every 30 seconds. However, on very long runs, this leakage uses up high memory, which can indirectly impact other running daemons. PR1324531
- On 30-Port MacSec-enabled line card (LC1101-M-30C, LC1101-M-30Q, and LC1101-M-96X) of the PTX10000 chassis, when the exclude-protocol lacp statement configured at the [edit security macsec connectivity-association connectivity-association-name] hierarchy level is deleted or deactivated, the LACP protocol's Mux State shown under the output of CLI command show lacp interface, might remain as attached or detached and might not change to distributing state. PR1331412
- With regards to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infra relies on the integrity of the TCP connections. The reactions to failure situations might not be handled gracefully, resulting in TCP connection timeouts because of jlock hog crossing the boundary value (5 seconds), which causes bad consequences in MX Series Virtual Chassis. Currently, there is no other easy solution to reduce this jlock hog besides enabling marker infra in the MX Series Virtual Chassis setup. PR1332765
- 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
- On MX2010 and MX2020 routers equipped with SFB2, some error logs might be seen. PR1363587
- On the EX9208 device, few xe interfaces go down with the error message **if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error**. PR1377840

- 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
- The virtio throughput remains the same for the multi-queue and single-queue deployments. PR1389338
- On MX platform with enhanced-ip and VRRP configured, if you remove or add a child link from aggregated Ethernet bundles, traffic destined to VRRP VIP might be dropped. PR1390367
- Revert of RLT to primary might silently discard traffic for around 10 minutes after the primary FPC is online with primary RLT up.PR1394026
- The FPC generates core files under certain circumstances on addition and deletion of hierarchical CoS from pseudowire devices.PR1414969
- Traffic statistics are not displayed for the hybrid access gateway session and tunnel traffic. PR1419529
- If the HTTP header enrichment function is used, the traffic throughput decreases when the traffic passes through header enrichment. PR1420894
- When you run the **show route label X | display json** command, two **nh** keys are present in the output. PR1424930
- Dynamic tunnel summary displays an incorrect count of up and total tunnels after multiple iterations of activating and deactivating the dynamic tunnel configuration. It is just a display issue and there is no problem with the functionality. PR1429949
- Customers need to know to which shard a route will hash (even before route coming up). The below new test command can show it. Test rib-sharding route <destination>. Prerequisite: Sharding should be configured. Below is sample sharding configuration. set system processes routing force-64-bit set system processes routing bgp rib-sharding number-of-shards 4 set system processes routing bgp update-threading number-of-threads 1 Below are the examples of test command when 4 shards are configured. root@r1_re> test rib-sharding route 1.1.1.1/32 Prefix: 1.1.1.1 hash to junos-bgpshard1 root@r1_re> test rib-sharding route 1.1.1.4/32 Prefix: 1.1.1.4 hash to junos-bgpshard0 PR1430460
- On MX Series platforms, if the clock frequency is slowly changing on CB0 (slow drift), the clock source for MPC-3D-16XGE-SFPP may not be changed to CB1, which will cause interfaces on it to go down and remain in the down state. PR1433948
- On the dual Routing Engines of the MX Series platforms with subscriber management, the replication daemon (repd process) might crash after booting for the first time with a newly installed Junos OS release. The repd process synchronizes subscriber information across Routing Engines, so normally the repd crash has no impact on the live service. PR1434363
- On MPC10E 3D MRATE-15xQSFPP, Layer 2 over GRE is not supported. Although the configuration gets committed, the feature does not work. PR1435855
- Two extra fields - sni and url-category - are added to SSL_PROXY_INFO logs. This information helps identify the SNI information received and which UTM category that falls under, when a session is whitelisted. The field sni is populated with the value from the CLIENT HELLO extension "server name

indication"; and the field url-category is populated with the UTM category the sni info in CLIENT HELLO matches. In case, no UTM category is matched url-category value is NA. PR1442391

- Interface hold-down timers cannot be achieved for less than 15 seconds on the MPC11E line card.PR1444516
- In corner cases following information will be observed on FPC: Jul 21 10:47:13 [TRACE] [router] [Jul 21 04:52:10.294 LOG: Err] mqss_pio_read_u32: Reading 32-bit register failed - status 1, pio_handle 0xe1b445c0, addr 0x4001150 Jul 21 10:47:13 [TRACE] [router] [Jul 21 04:52:10.294 LOG: Err] mqss_ptc_intr_handler: Reading 'enable' failed - status 1, ptc_block_num 0 Jul 21 10:47:13 [TRACE] [router] [Jul 21 04:52:10.294 LOG: Err] mqss_ea_host_misc_wanio_intr_handler: Processing the PTC interrupts failed - status 1 Jul 21 10:47:13 [TRACE] [router] [Jul 21 04:52:10.294 LOG: Err] mqss_host_misc_wanio_intr_handler: Processing WANIO interrupts for EA failed - status 1 Jul 21 10:47:13 [TRACE] [router] [Jul 21 04:52:10.294 LOG: Err] mqss_host_misc_wanio_intr_callback: Processing the WANIO interrupts failed - status 1 Jul 21 10:47:13 [TRACE] [router] [Jul 21 04:52:10.300 LOG: Err] mqss_pio_read_u32: Reading 32-bit register failed - status 1, pio_handle 0xe1b445c0, addr 0x4001150 Jul 21 10:47:13 [TRACE] [router] [Jul 21 04:52:10.300 LOG: Err] mqss_ptc_intr_handler: Reading 'enable' failed - status 1, ptc_block_num 0 Jul 21 10:47:13 [TRACE] [router] [Jul 21 04:52:10.300 LOG: Err] mqss_ea_host_misc_wanio_intr_handler: Processing the PTC interrupts failed - status 1. PR1450358
- Physical interface policers are not supported in Junos OS Release 19.3 for the MPC11 line card.PR1452963
- Cosmetic issue that affects only CLI. RADIUS, L2TP, and so on are unaffected. The CLI issue is seen after ANCP restart and before ANCP neighbor is re-established and port-ups are received. Under normal working conditions, after ANCP restart, the port-ups should be received right away and the CLI issue will not be seen. PR1453837
- With logical system configuration, filter-based GRE encapsulation does not work. PR1456762
- With the scaled filter-based forwarding (FBF) configuration, two instances seem unable to forward the traffic to the respective routing instances. It appears that the FBF programming is incorrect for the two FBF instances. PR1459340
- Occasional warning messages such as **TCP connect error** can be seen during the FPC reboot. These are generally inconsequential and have no impact on the FPC or the line-card software functionality.PR1460153
- The CFM remote MEP is not coming up after configuration or remains in start state. PR1460555
- On MX platforms with MPC10E/MPC11E, a BFD session might flap when it is moving to an aggressive interval after coming up with a slow/nonaggressive interval. PR1462775
- Backport jemalloc profiling CLI support to all releases where jemalloc is present. PR1463368
- The traffic stops when volume quota is reached but resumed wrongly after APFE failover. Threshold and quota values are not updated to the secondary APFE. If quota is hit on the primary APFE and traffic starts dropping due to quota and switchover happens, traffic will continue to flow until quota is hit.PR1463723

- Harmless syslog error messages might be seen and expected during ISSU or GRES or FPC offline/online scenarios. PR1464524
- A BFD session might flap when it is moving to an aggressive interval after coming up with a slow/nonaggressive interval. This issue is mainly seen in a scaled setup. PR1465285
- On Junos device, if a member link of aggregated Ethernet is copper cable and the speed is set to 100mbps in remote end explicitly or through auto-negotiation, aggregated Ethernet interface might go down after GRES if the aggregated Ethernet speed is different than 100mbps. PR1465593
- On all MX platforms, if a GRE (Generic Routing Encapsulation) tunnel destination is configured to be part of a routing instance, traffic with "IP/AH (Authentication Header)/GRE" headers might be forwarded incorrectly. These kinds of headers can be seen in IPSec transport mode. Headers like "IP/AH/IP/GRE" can work well in most cases. PR1466062
- Unable to get the service sessions when NAT64 is configured with the destination-prefix length 32. PR1468058
- The MPC11E line card might take additional time to come during the movement from one GNF to another GNF in an in-chassis junos node slicing setup. PR1469729
- With BGP rib-sharding and update-threading, traffic drops 100 percent in the BGP Layer 3 VPN streams, after the removal or restoration of the configuration. PR1469873
- The following message might be seen in chassisd log after rebooting or changing the configuration, etc.: re_tvp_builtin_fwinfo_update: Unable to get firmware version. PR1471938
- Syslog error message **Failed to complete DFE tuning** is generated. This message has no functional impact and can be ignored. PR1473280
- For the MPC10E line card, the IS-IS and micro-BFD sessions do not come up during baseline.PR1474146
- When the external server is rebooted, the SNMP values configured within the /etc/snmp/snmpd.conf file at the server gets overwritten with the content from the JDM SNMP configuration section. The trap configuration changes get completely removed. Restarting or stopping and starting JDM does not change the host /etc/snmp/snmpd.conf file. PR1474349
- Error message **[Error] L2alm : l2alm_mac_process_hal_delete_msg:667 Ignoring MAC delete with ifl index 355, fwd_entry has 7888** are seen after performing configuration removal/restoration with IP/MPLS configurations in the MX480. PR1475785
- In VPLS configurations, ARP resolution over an IRB interface might fail if the hosts are behind a vt-tunnel. As a workaround, you can use **no-tunnel-services** statement. PR1477005
- On the MX platforms with the type of 3D 20x 1GE MIC installed, after performing ISSU (In-Service Software Upgrade), the FPC equipping the MIC may crash and interfaces stay down. Due to this issue, the traffic on the MIC will be impacted. PR1480212
- Invalid packets are dropped by DUT with TCC encapsulation configuration as intended but the statistics counters are incremented. PR1481698

- When an SFB3 is un-gracefully yanked out and plugged back in and comes back online, there is a chance that the SFB3 might end up in error state during fabric chip initialization error. This error is seen intermittently. PR1482000
- The next-generation services MX Series SPC3 services card can exhibit inconsistent behavior when the vmhost image is installed on the next-generation Routing Engine (NG-RE) RE-S-X6-64G-UB. Other Routing Engines that are compatible with next-generation services do not experience this problem. These Routing Engines are RE-S-1800X4-16G-UPG-BB, RE-S-1800X4-32G-UB, RE-S-1800X4-16G-UPG-BB, and RE-S-1800X4-32G-UB. PR1482334
- Possible out of order deletion of AftNode #012#012#012 AftNode details - AftIndirect token:230791 group:0 nodeMask:0xffffffffffffffff indirect:333988 hwInstall:1#012 "during baseline PR1486158
- Failed to open session database and unexpected LSP flap is observed on ingress node after ZPL ISSU. As a workaround, use the **set protocols lacp fast-hello-issu** command. PR1488089
- Incase of G.8273.2 measurement of asymmetric combination of port speeds involving 40G , we observed cTE metrics of 60-120ns in certain trials. PR1488549
- Next-hop learning command is enabled by default in the MPC10 and MPC11 line cards irrespective of the command configuration. PR1489121
- The **show ptp statistics detail** command shows incorrect values for the delay request and response packets.PR1489711
- Accessing the free memory might fail even after multiple switchovers of more than 50 with scale configuration generating a core file. PR1491527
- On MX204 and MX10003 routers with the MPC7E, MPC8E, MPC9E, MPC10E, and JNP10K-LC2101 line cards, the following error message is observed: **unable to set line-side lane config (err 30)**. This is error does not have any impact and can be ignored.PR1492162
- The component sensor does not export data under CBO/1 in the expected time. PR1493579
- In a VPLS scenario after an NSR Routing Engine switchover, the flood next-hop ID for the VPLS instance might not get synchronized between primary and backup Routing Engines which might lead to the traffic loss for that VPLS instance.PR1495925
- After backup Routing Engine halt, CB1 goes offline and comes back online; this leads to the backup Routing Engine booting up and the reboot reason is displayed as **0x1:power cycle/failure**. There is no functional impact of this issue.. PR1497592
- In routing-instance with table next-hop scenario (e.g. if EVPN routing-instance is configured, the l2ald process creates a routing table and the EVPN adds a route pointing to this table as table next-hop in the rpd process), if the routing table created within the routing-instance is deleted and then re-added (e.g. deactivated and then re-activated the routing-instance) very fast before the rpd could delete the route pointing to the table next-hop, then the route in the rpd will end up using the staled table next-hop, hence resulting in traffic loss. Sampling configuration which delays the route deleting in the rpd increases the possibility of hitting the issue. PR1498087

- On the certain MX platform with MACsec configured, the packets < 64B over MACsec are getting dropped due to the MACse chip on the MPC will discard the packets as runt packets by default. PR1499966
- On an MX2020 device, the SFB3 and MPC11 line card is not supported. This change is to disable these components in JUNOS software version 19.4. PR1503605
- On MX204 and MX10000 series or platforms with MPC7/8/9/10/11, in many cases, other CMERRORs will be invoked as well and Major Alarms will perform disable-pfe action. However, in some cases of FO/WO errors, this does not happen. The fix is to create a new CMERROR if the Packet Error count of 255 is active for 3 consecutive polling periods to cover the condition with operational impact but no other CMERROR events have caught this event. PR1503705
- The output of Aggregate Ethernet interface statistic does not include its member links' statistics. PR1505596
- In an EVPN scenario with VRRPv6, the Ethernet source MAC address might be used for IPv6 MAC-IP binding when the NA is sent from the VRRPv6 primary. AS this unexpected behavior is triggered on regular intervals, it causes the entries to keep refreshing in the EVPN database because NS from the VRRPv6 primary changes the MAC-IP binding. This impacts the traffic PR1505976
- VLAN membership for interfaces is not added when configured with the vlan name of a VLAN created using vlan-id-list.PR1506045
- When a route resolves over a comp NH and its target NH is also resolves over another comp NH, the order of labels pushed might not be correct. PR1508644
- The WAN-PHY interface continuously flaps with the default hold-time down of value 0. PR1508794
- In Junos Node Slicing with the MPC11E scenario, the MPC11E is introduced as part of the GNF (guest network function) setup. If the link between the external server and CB (control board) goes down, the MPC11E might not boot fully but remain in the "READY" state. PR1510358
- BGP-SRTE binding-sid with more than one label stack Need's enhancement for PTX10003-80c PTX10003-160c. For PTX does not support more than 1 chain-composite-next-hop, it can configure this cli knob to disable chain-composite-next-hop as work around. with this knob configured, all labels will pushed from egress. protocols { source-packet-routing { no-chained-composite-next-hop; } } PR1512213
- On MX2000 platforms with MPC11, during system reboot or MPC11 linecard reboot, the impacted MPC11 linecard might not be able to boot up normally due to this timing issue, and it might take extra 15-30 seconds to boot up again in the worst situation. It is a rare issue. PR1514090
- A delay of 35 seconds is added in reboot time in Junos OS Release 20.2R1 compared to Junos OS Release 19.4R2. PR1514364
- In TCP log scenario, the TCP session to syslog server might be down. This issue has service impact. PR1519896
- After doing NSSU and performing GRES backup member of Virtual Chassis might go to unstable state leading to kernel core with db>prompt. This issue can impact traffic. PR1533874

- Inconsistent core.python2.7.mpc0 core files are seen.PR1534568
- On MX2010 and MX2020 platforms running in Junos Node Slicing scenario, when the base system (BSYS) has Routing Engine (RE) switchover which is followed by the addition/deletion of MPC11 to/from the guest network function (GNF), if the slot number in use for MPC11 is 8 or above, the MPC11 linecard might be stuck in ready state due to this issue. PR1535588
- In scaled MX2020 router, with vrf localisation enabled, 4 million next-hop scale, 800,000 route scale. FPCs may go offline on GRES. Post GRES, router continues to report many fabric related CM_ALARMs. FPC may continue to reboot and not come online. Rebooting master and backup RE will help recover and get router back into stable state. PR1539305
- RPD crashes in rt_nh_cache_entry_delete() during initial config or steady state. BGP multipath is present in config and route scale is high. Exact trigger or series of events that cause the crash is unknown. PR1544570
- Even though enhanced-ip is active, the following alarm is observed during ISSU: **RE0 network-service mode mismatch between configuration and kernel setting**. PR1546002
- 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 PR1548677
- In rare cases of power related failures on the FPC, Fabric Healing will detect and try to heal this fault condition by performing an offline or online FPC event. If the same FPC fails again within 10 minute period, fabric auto-healing attempt is considered failing and the FPC will get off-lined to avoid further operational impact. If during the power offline event, the faulty FPC gets disconnected ungracefully due to the hardware power fault, the FPC might attempt an on-lined request again after 5 minutes. There may be traffic impact due to this issue. PR1556558
- The MAC addresses may fail aging out under an environment Virtual chassis where a large number of MAC addresses are learned. This issue was observed with mac entries 280000 in the Virtual chassis devices. PR1558128
- This feature was never tested or claimed to be working on this plaftform(pvi-model/MX10000 series). The feature does not need any extra effort (or atleast from the feature side) to support on the new platform. So, it is believed that if it is supported, we can create a test-only RLI and track via the same. PR1559200
- On all Junos platforms with MPC10E or MPC11E linecards, repeated link flaps on an interface could result in complete traffic stall (packets no longer going out of the interface). PR1560772
- On all L2NG platforms, MAC address entries might be smaller in the MAC table than in the ARP table, this because some of MAC addresses are not relearned successfully after MAC address age timeout. This issue will cause traffic loss for non-existing MAC entries. PR1567723
- Traffic might be dropped on MX trinity platforms when the default route is changed in the inet.0 table. It might take 2-3 seconds to be updated in Packet Forwarding Engine. This issue can be recovered automatically. PR1568944

- On the MX Series platforms, the jinsightd process might be stuck with high CPU process utilization if the services jinsightd is enabled in the Junos Telemetry Interface (JTI) scenario. PR1570526
- On MX platforms, if sample based action is used in firewall filter for an interface, such as syslog/log/port-mirror/Jflow, traffic loss might be observed if the sampled packets rate exceeds default DDOS sample bandwidth. PR1571399
- With the enabled 'persist-groups-inheritance' configuration statement "Licenses used" fields in 'show system license' output are showing up all zeros despite the fact that the necessary licenses were installed and LNS/AAA configuration applied. PR1572507
- The following messages might be seen in the logs from MPC11E line-card: Feb 9 11:35:27.357 router-re0-fpc8 aftd-trio[18040]: [Warn] AM : IPC handling - No handler found for type:27 subtype:9. There is no functional impact, these logs can be ignored. PR1573972
- The PIM rib-group failure to add in vrf - PIM: ribgroup vrf not usable in this context; all RIBs are not in instance. PR1574497
- **CHASSISD_FRU_IPC_WRITE_ERROR: fru_send_msg: FRU GNF 2, errno 40, Message too long** might appear periodically in the chassisd logs. PR1576173
- The analyzer may not work properly when port channelization is done after the analyzer configuration and with these newly created ports new analyzer is created. PR1580473
- The problem happens due to improper cleanup of reference count, hold on interface. Reference count are added on logical interfaces, so dependencies gets cleaned up properly. We could see **DCD_CONFIG_WRITE_FAILED** error messages flooding in /var/log/messages when problem gets hit. PR1581260
- On MX104 platforms with 3D 20x 1GE(LAN) SFP MIC installed, if the 100M SFP (100Base-LX SFP) is inserted to the MIC, when the "hierarchical-scheduler" or "per-unit-scheduler" is configured or removed from the interface of the 100Base-LX SFP, the interface will stop transporting traffic correctly. PR1582724
- On MX platforms running Junos, traffic might not get filtered as per whitelist and blacklist. The issue happens while deactivating and activating the security-intelligence profile when web-filtering is configured on the box. During this, the whitelist and blacklist files from PE are not being added to the Packet Forwarding Engine filter causing the issue. PR1584377
- MPC2E NG PQ & Flex Q with MACsec dropping 50 percent of traffic after enabling flexible-queuing-mode on both 1G and 10G.Changed the port speed for MACSEC MIC from 1G to 10G to increase the XQIF queue size to 32. On MX platform with MPC2E NG line cards and MACSEC MIC, traffic drop of about 50 percent may be seen when flexible-queueing-mode is enabled and a traffic is sent at the rate of more than 5Gbps. PR1586403
- In the BBE based with ACI (Agent Circuit Identifier) VLAN Interface Sets scenario, the ACI VLAN ifl/iflset pointer should be freed properly after performing the clean-up operation of the subscriber sessions (e.g. logging out the subscribers, and so on). But, in some corner cases, if the ACI VLAN ifl/iflset pointer gets freed well before performing the clean-up operation, it might become the invalid one in the system.

Then the bbe-smgd might crash since the subscriber sessions try to access this kind of invalid pointer. PR1587792

- On all Junos and EVO platforms, when there is a congestion on the link where telemetry streams are connected, then in a race conditions, there can be na-grpcd core and telemetry service will be impacted as na-grpcd will take a minute to come back online. PR1587956
- On all Junos platforms, in a rare scenario with scaled routing set up, the Kernel memory might get full which could lead to the rpd crash. There will be service impact and it will get recover automatically after the crash. When the rpd crash happens, the core-dump files could be seen by executing CLI command **show system core-dumps**. user@hostname>show system core-dumps -rw-rw- - - - 1 root field /var/tmp/rpd.core<*>.gz PR1588439
- When PCP mappings timeout, then NAT port release syslog messages are not getting generated. This issue is seen in MX-SPC3. PR1591297
- On MX platforms with Next Generation Subscriber Management (Tomcat) enabled, if the COS CR-features (Classifier/Rewrite/Frag-map) are used by the VBF (Variable Based Flow) service, the MPC might crash in a rare case. The specific trigger is not known as this issue cannot be able to replicate. PR1591533
- On the platforms with MS-MPC/MS-MIC used, it can be seen that the "clear-ipsec-sas-for-duplicate-ts" can't clear Secure Access (SA) for duplicate traffic-selectors (TS) if the local traffic selector (TS) is different and the remote traffic selector (TS) is same. PR1591735
- In a rare case, the logical interface (IFL) of aggregated interface (e.g., AE, RLT, RVT, AF, AMS, RLSQ interface etc.) might fail to be added to Junos kernel. In this case, the Routing Engine kernel might crash with vmcore file generated. The IFL of aggregated interface adding failure in Junos kernel could happen in cases like failure of multicast filter list initialization or DCD sending an invalid vlan-id or memory allocation error etc. PR1592456
- On MX platforms with SPC3 used, if adding the PS interfaces on the Routing Engine after SPC3 is up and running, the packet from the PS interface and is sent to SPC3 for services like NAT/SFW/IDS, etc. might be dropped by SPC3. PR1592706
- On MX platforms with MS-MPC and MS-MIC when tcp-tickle knob is enabled under services-options in DS-lite (Dual-Stack lite) with NAT scenario, the TCP keepalive might not be processed by the private network host and the purpose of TCP keepalive gets compromised. PR1593226
- The MX5, MX40, and MX80 TEB may stuck in Present state after upgrade or chassis reboot. This is because the Routing Engine mastership setting fails during the chassis boot up. In the fixed Junos code, one extra check has beed added to check the readiness of mastership setting and a retry mechamsim is also added to avoid such scenario. PR1595107
- In Enhanced Subscriber Management environment with interim-update configured, if a subscriber is present over multiple Packet Forwarding Engine instances (e.g. configure subscriber interface over aggregated Ethernet bundle), which is hosting in Push-model MPC that supports Next-Gen Broadband-Edge Statistics (e.g. MPC2E-NG/MPC3E-NG, MPC5E/7E/8E/9E), the wrong Input/Output Octets and Packets count in Interim-update may be observed. Please note that this issue is only applicable

for releases 20.1 and prior. The code was restructured on 20.2 and post, so the issue is no longer applicable. PR1596645

- On MX10003 platform with MACsec used scenario, traffic loss might happen periodically if Routing Engine is working under a pressure situation (rpd memory occupied around larger than 70%), which may cause the message of Secure Association Key (SAK) of MACsec to be vetoed by kernel that causes one of pair (RX/TX) Secure Association (SA) number missing. Moreover, the missing SA number is still available in the system, so whenever SA number is rollover to it (SA number is rollover between 0 to 3), traffic loss might happen due to invalid SA pair. PR1596755

Class of Service (CoS)

- CoS EXP classifier and rewrite with the **mpls-inet-both-non-vpn** protocol option is not working as expected. PR1479575
- On all Junos platforms with LT/PS interface is configured, default classifier in wildcard will get attached to the LT/PS interface, even if no classifier is configured. It could be observed when there is a wildcard interface which is matching the LT (Logical-Tunnel)/PS (Pseudowire Subscriber) interface. PR1542559

EVPN

- In an EVPN scenario, duplicate packets are seen because a nondesignated forwarder is sending an inclusive multicast packet to the PE-CE interface after MAC lookup. PR1245316
- In an EVPN scenario with nonstop active routing (NSR) enabled, the rpd crashes and generates core files on the backup Routing Engine when any configuration changes on the primary Routing Engine.PR1336881
- With Junos OS Release 19.3R1, the VXLAN OAM host-bound packets are not throttled with DDoS policers. PR1435228
- On all Junos platforms in EVPN-VXLAN to EVPN-MPLS stitching scenario, traffic loss could be seen with DF (Data Forwarder) changes when traffic flows from VxLAN to MPLS. The traffic loss will occur till mac-ip ages out. PR1515096
- In PBB-EVPN (Provider Backbone Bridging - Ethernet VPN) environment, ARP suppression feature which is not supported by PBB might be enabled unexpectedly. This could cause MAC addresses of remote CEs not to be learned and hence traffic loss. PR1529940
- On all platforms that support EVPN with IGMP Snooping scenario, multicast traffic might be dropped at receiver PE when the EVI-RT community value carried in Type-7 route is a 4-byte AS number instead of 2-byte and it results in out-of-sync status for Type-7/8 route. PR1582134

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.

Forwarding and Sampling

- Packet length for ICMPv6 is shown as "0" in the output of the **show firewall log detail**CLI command. PR1184624
- When an IPv4 prefix is added to a prefix list referenced by an IPv6 firewall filter, the log message **Prefix-List [Block-Host] in Filter [Protect_V6] not having any relevant prefixes , Match [from prefix-list Block-Host] might be optimized** is not seen.PR1395923
- When GRES is triggered by SSD hardware failure, the syslog error **rpd[2191]: krt_flow_dfwd_open,8073: Failed connecting to DFWD, error checking reply - Operation timed out** might be seen. Issue can be recovered by restarting the dfwd daemon. PR1397171
- After you restart routing, the remote mask, which indicates from which remote PE devices MAC IP addresses are learned, that the routing daemon sends might be different from the existing remote mask that the Layer 2 learning daemon had before restart. This causes a mismatch between Layer 2 learning and the routing daemon's interpretation as to where the MAC IP address entries are learned, either local or remote, leading to the MAP IP table being out of synchronization. PR1452990
- Expected number of 512,000 MAC addresses are not re-learned in the bridge table after clearing 512,000 MAC addresses from the table. PR1475205
- The IPv6 filter might not capture the host outbound traffic with the expected forwarding-class match condition. This can result in IPv6 host outbound traffic forwarding control issues. PR1491492
- After applying a list of firewall filters via 'input-list' or 'output-list' statement (or both) within the 'filter' stanza on a logical interface (IFL), the resulting concatenated filter is interface-specific. The system-generated name of the interface-specific filter consists of the full interface name followed by either '-i' for an input filter list or '-o' for an output filter list (like ge-1/3/0.0-i or ge-1/3/0.0-o). The system-generated name of the firewall filter counter consists of the configured counter name followed by a hyphen ('-') and this resulting concatenated filter name (like counter1-ge-1/3/0.0-i or counter1-ge-1/3/0.0-o). The firewall filter policer uses similar naming rule as firewall filter counter. If an IFL is configured with a firewall filter list for different families, the name of the resulting concatenated filter for different families will be same. This will not cause any issue on CLI. However, if the firewall filter telemetry data is streamed via Junos Telemetry Interface (JTI), it might cause confusion on collector side because the firewall filter list for different families will be treated as one filter. In particular, if firewall filters having same firewall filter counter (or policer) name are used in firewall filter list for different families, the incorrect statistics might be seen on collector because the firewall filter counter (or policer) name for different families cannot be distinguished on collector side. In order to avoid the issue, this PR

fix changed the name of the resulting concatenated filter for firewall filter list and added family to the filter name. The new system-generated name of the concatenated filter consists of the full interface name followed by a hyphen ('-') and family, and then either '-i' for an input filter list or '-o' for an output filter list (like ge-1/3/0.0-inet-i or ge-1/3/0.0-inet-o). PR1514141

High Availability (HA) and Resiliency

- When ZPL is done while traffic is running, some BGP sessions might flap leading to traffic loss. The drop is transient and traffic recovers post ZPL. PR1487144

Infrastructure

- On MX platform the harmless log of **invalid SMART checksum**might be seen when performing software upgrade to specific releases (for example, Junos OS Releases 15.1F5-S3, 15.1F6-S1, 15.1F7, 15.1R4-S3, 15.1R5, 16.1R1, 16.1R2, 16.2R1). PR1222105
- 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. PR1232403
- The following error message is seen during FTP: **ftpd[14105]: bl_init: connect failed for /var/run/blacklistd.sock(No such file or directory)**.PR1315605
- F-label veto code checks for per-pfe f-label pools. PR1466071
- **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
- User while loading the kernel would see the message **GEOM: mmcsd0s.enh: corrupt or invalid GPT detected**. This message has no impact to functionality and will be resolved in a future release. PR1549754
- Memory corruption of a binary from /usr/bin/ or /usr/sbin/ directory can occur if such binary is invoked when a recovery snapshot creation is in progress. The exact symptoms are different depending on the exact binary and JUNOS version. Some programs show an error, and some programs crash every time it is executed. Such memory corruption is persistent until the affected Routing Engine is restarted. Please refer to TSB17954 (https://kb.juniper.net/TSB17954) for further details. PR1563647

Interfaces and Chassis

- In an L2TP scenario when an MX Series router functions as an LTS (L2TP tunnel switch), there is a memory leak in the jpppd process running on the backup Routing Engine. This eventually generates jpppd core files due to an out-of-memory condition. There is no functional impact as it happens on the backup Routing Engine. PR1350563
- The SFP index in the Packet Forwarding Engine starts at 1, while the port numbering starts at 0. This causes confusion in the log analysis. PR1412040
- Layer 2 logical interface configuration is now decoupled from bridge or EVPN configuration. A Layer 2 logical interface can now be configured without being assigned to a bridge/EVPN. PR1438172
- When a user checks for interface-level statistics, an issue occurs for IPv6 counters. At the originating router IPv6 local statistics counters are not updating because IPv6 local statistics counters not incrementing. IPv6 transit statistics are derived from total statistics and local statistics (Transit statistics = Total - Local). Because the local statistics are not updating, total statistics and transit statistics will be the same. This issue is specific to platforms with MPC10E and MPC11E. PR1467236
- Changing framing modes on a CHE1T1 MIC between E1 and T1 on an MPC3E NG HQoS line card causes the PIC to go offline. PR1474449
- The traffic (which is destined to the hosts behind static PPPoE subscriber's CPE device) drop is seen due to bad MPLS VPN label (which points to discard next-hop) after Routing Engine switchover without NSR. The traffic destined to the CPE device itself is not affected. PR1488302
- New SCB cards may have uninitialized VC Data Blocks, preventing setting the member-id when configuring as a MX-VC for the first time. PR1569556

J-Web

- An Improper Input Validation vulnerability in J-Web of Juniper Networks Junos OS allows a locally authenticated attacker to escalate their privileges to root over the target device. Please refer to https://kb.juniper.net/JSA11182 for more information. PR1592021

Junos Fusion Provider Edge

- On Junos fusion system, intermediate traffic drop is sometimes seen between AD and SD when sFlow is enabled on the ingress interface. When sFlow technology is enabled, the original packet is getting corrupted for those packets that hit the sFlow filter This due to few packets transmitted from the egress of AD1 is short of FCS (4 bytes) + 2 bytes of data due to which the drops occur. it is seen that the normal data packets are of size 128 bytes while the corrupted packet is 122 byte. PR1450373

Layer 2 Features

- Traffic does not get load balanced over ESI links with EVPN_VXLAN configured. PR1551543

Layer 2 Ethernet Services

- Only hitless CB upgrade is supported and not CB downgrade. So if something happens during or after the SCB upgrade and the customer want to revert back from E-SCB to SCB, a system restart is mandatory. PR980340
- The DHCP DECLINE packets are not forwarded to the DHCP server when **forward-only** is set within **dhcp-reply**.PR1429456
- On MX platforms with DHCP ALQ, the ALQ (Active Lease Query) TCP Queue may get stuck. This might cause the subscribers from Backup BNG (Broadband Network Gateway) not to be able to sync with primary BNG and eventually causing the subscribers in primary to start going down and result in a major outage. PR1590421

MPLS

- The rpd generates core files at **hbt_iterate_next, ldp_purge_unknown_tlv_temp_tree**. PR1210526
- On the MPC11E line card, degradation in IS-ISv4/RSVP convergence with link-protection configuration is observed. PR1485701
- When ISIS-TE or OSPF-TE is enabled without admin-groups-extended-range/admin-groups-extended (which is configured under routing-options) or admin-group-extended configured, it receives the peer-router advertised the extended admin groups and then enable the config of

admin-groups-extended-range/admin-groups-extended and admin-group-extended, some LSP with extended admin group constraints will fail to be established. PR1575060

Network Management and Monitoring

Platform and Infrastructure

- Packet header might get corrupted at the ingress Packet Forwarding Engine if the packet with more than two IEEE 802.1Q VLAN tags are traversing in an EVPN/VPLS routing instance. PR1300211
- On MX-Series platforms with MPC7/8/9 or MX-204/MX-10003 when the packets which exceed the MTU and whose DF-bit is set go into a tunnel (such as GRE, LT), they might be dropped in the tunnel egress queue. PR1386350
- On MX Series routers with MPCs, the unicast traffic might drop when the destination is reachable over an integrated routing and bridging (IRB) interface and a label-switched interface (LSI) with two next hops. PR1420626
- The traps are the result of PPE commands injected from the host. One possible reason could 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 ISSU and this may be due to a problem with ISSU counter morphing used for LU-based cards, where certain counters are not disabled or disabled too late during ISSU. PR1426438
- On the MX480 devices, traffic loss is observed if the ingress and egress ports are on different FPCs. PR1429714
- For the bridge domains configured under an EVPN instance, ARP suppression is enabled by default. This enables the EVPN to proxy the ARP, and reduces the flooding of ARP in the EVPN networks. Because of that, storm control does not take effect on ARP packets on the ports under such bridge domains. PR1438326
- A dual Routing Engine Junos node slicing GNF with no GRES configured and with **system internet-options no-tcp-reset drop-all-tcp** configuration might enter dual backup Routing Engine state upon manual GNF Routing Engine primary-role switchover attempt with the **chassis routing-engine master [acquire|release|switch]** command from either GNF Routing Engine CLI. PR1456565
- While SNMP-Agent polls round-trip time (RTT) related to OIDs from a router running Junos OS, such as pingResultsAverageRtt, the router might respond with zero (0) value even there is no RPM ping failure. The following objects might be impacted: **iso.3.6.1.2.1.80.1.3.1.4 -> pingResultsMinRtt iso.3.6.1.2.1.80.1.3.1.5 -> pingResultsMaxRtt iso.3.6.1.2.1.80.1.3.1.6 -> pingResultsAverageRtt iso.3.6.1.2.1.80.1.3.1.7 -> pingResultsProbeResponses iso.3.6.1.2.1.80.1.3.1.9 -> pingResultsRttSumOfSquares.**PR1458983
- After performing ISSU with a scaled configuration, high CPU utilization is observed in the MPC 3D 16x 10GE card. PR1461715

- With multiple different fixed-sized traffic streams configured at 1,000,000 fps (40 Gbps combined rate) on aggregated Ethernet0 along with another independent aggregated Ethernet interface (aggregated Ethernet1, 50 percent line rate 4 streams bidirectional => 118 Gbps combined traffic rate), both hosted on a single Packet Forwarding Engine instruction of an MPC11E line card, small varying packet drops occurs for every iteration on aggregated Ethernet1 on disabling aggregated Ethernet0. The drops might vary from 200 to certain 1000 frames. PR1464549
- Line-card errors found at **RT-HAL,rt_mesh_group_delete_check,1599: Deletion of a non-existant mesh-group : proto 35 rtt 60grp-index 0,PFE_ERROR_NOT_FOUND: route check failed, entry not found with steady state measurement check**. PR1472454
- Unknown unicast filter applied in an EVPN routing instance blocks unexpected traffic. PR1472511
- A few OAM sessions are not established with scaled EVPN ETREE and CFM configurations. PR1478875
- An EVPN does not support individual logical interfaces operations if the ESI is configured on the physical interface. On MX Series routers, the loop prevention feature supports per logical interface flap, only when ESI is configured at the logical interface level. If one logical interface is flapped during BUM traffic flow through another logical interface, the loop prevention feature does not work. If ESI is configured on the physical interface and one of the ESI logical interfaces is down, EVPN considers it as whole physical interface down. The router's designated forwarder (DF) role or MH status is changed and the local bias filter is reconfigured or deleted. Because the traffic flowing through the logical interface is not flapped, loop prevention is not enabled for the logical interface.PR1485100
- On all platforms running Junos OS that support EVPN-MPLS and EVPN-VXLAN, when an existing ESI interface flaps or is added newly to the configuration, sometimes designated forwarder (DF) election happens before the local bias feature is enabled and during this time, existing broadcast, unknown unicast, and multicast (BUM) traffic might be looped for a short time duration (less than several seconds). PR1493650
- Traffic loss is observed after unified ISSU, when you enable or disable and activate or deactivate the interface. PR1493723
- With GRES and NSR functionality with VXLAN feature, the convergence time may be slightly higher than expected for L2-DOMAIN-TO-L3VXLAN PR1520626
- On the MX series platforms, a CFM session may not come up if it's configured along with encapsulation "vlan-ccc" on a logical interface. PR1522370
- On MX and EX9200 serial platforms, under Ethernet VPN (EVPN) environment, packets routed using IRB interface could not be fragmented due to media maximum transmission unit (MTU) problem. PR1522896
- 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

- Due to a rare problem in TCP socket replication between routing engines and md5 digest processing on the backup Routing Engine in NSR configuration, the new master Routing Engine may crash during Routing Engine switchover. PR1527246
- On the MX platforms with subscriber management service configured, if the distributed IGMP processing is enabled for subscribers, the MPC may crash occasionally due to the flaw of which IGMP and pfeman threads access IGMP related data structure. PR1534542
- On all Junos platforms with certain linecards(such as MPC7/8/...12), when the multicast/flood route with multiple next hops in a large scale, the BUM(Broadcast, Unknown unicast, Multicast) and multicast traffic will be discard. PR1535063
- On the MX platforms with XM chipset based line card installed, when the line card experiences the CMERROR **XMCHIP_CMERROR_DDRIF_PROTECT_WR_RD_SRAM_RUNN_CHKSUM**, the disable-pfe action will be involved. This issue will cause the Packet Forwarding Engine to be disabled and traffic lost. PR1568072
- On all MX platforms, the L2TP tunnel will not work with filter-based encapsulation for the breakout interface. This issue is seen as the parsing logic in Packet Forwarding Engine for getting the tunnel parameters could not handle breakout interface scenarios. PR1568324
- This issue might be seen only in back to back GRES in about more than 40 to 50 iterations. No workaround available and FPC gets restarted. PR1579182
- The issue is due to output byte count not getting updated properly. The script logs shows that there is no packet loss, There is no functional impact and will be taken up in the upcoming releases. PR1579797

Routing Policy and Firewall Filters

- Routing policy actions failed to configure **neighbor-sets** and **tag-sets**. PR1491795

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 running Junos OS. PR993870
- 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, which 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 ldp-synchronization enabled for OSPF. user@host> show ospf interface ae100.0 extensive Interface State Area DR ID BDR

ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity. As per the current analysis, 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
- The **show version detail** command triggers the following severity error logs: **mcsnoopd: INFO: krt mode is 1" "JUNOS SYNC private vectors set"**.PR1315429
- SCP command with routing option (-JU) is not supported. PR1364825
- Ukern memory leak and fpc core crash might happen when device configured link-node protection with labeled-bgp. PR1366823
- Bfd session flaps during ISSU only in mpc7e card (Bfd sessions from other cards of DUT to peer routers did not flap during ISSU). Issue is not seen frequently. PR1453705
- Even when the **protocols mpls traffic-engineering bgp-igp**statement is configured, the UDP tunnel routes are not added to inet.0. The UDP tunnel routes are added only to inet.3 table whether the statement is configured or not. PR1457426
- With NSR enabled, the current BGP design supports 3000 BGP IPv6 peers or 8000 BGP IPv4 peers. If you try to bring up more than 3000 BGP IPv6 sessions or more than 8000 BGP IPv4 sessions, the rpd might crash.PR1461436
- In a BGP scenario, when the local router receives BGP updates with the same prefix from different BGP peers in a certain order with certain set of attributes, at this point a continuous soft-core may appear, but it allows the running rpd to continue. The act of taking a soft-core will put some CPU utilization for the time of taking the soft-core. PR1481933
- The virtual-router option is not supported under routing-instance in lean rpd image. PR1494029
- On all Junos platforms with scaling MVPN scenario, some PIM Join/Prune messages may not be processed for the first attempt. For instance, a dedicated PIM router receives more than 2500 PIM hello packets from the new neighbors, followed by PIM Join packets for the same multicast group in a very short period of time. PR1500125
- Disruptive switchover (no GRES or NSR configured) can lead to stale PPM entries programmed on the new master Routing Engine. If both GRES and NSR are activated after disruptive switchover and then a GRES switchover is performed, BFD sessions might flap continuously. PR1518106
- On the devices with NG-RE (Next Generation Routing Engine) and SCBE2 (Enhanced Switch Control Board), when BFD authentication for BGP is enabled, the BFD may flap after the NG-RE switchover. The switchover should be GRES or NSR switchover. After the flap, the device could be self recovery. PR1522261

- The BGP signaled dynamic tunnels remain in an established state and don't go down after deactivating the BGP export policy which contains next-hop and tunnel community information from the remote end. PR1579225
- In BGP multipath scenario, if an interface for a single hop EBGP peer goes down, the rpd might crash on the backup Routing Engine. If NSR switchover is performed, the rpd crash might be observed on the newly master Routing Engine, hence there may be traffic impact. PR1589141
- With OSPF remote LFA feature enabled, when ABR (area border router) with the primary interface and the secondary interface are in different OSPF areas, if the secondary interface is supposed to be chosen as part of the Remote-LFA path then the remote LFA backup path might not be formed. PR1592424

Services Applications

- In a subscriber management environment CLI commands **show services l2tp tunnel extensive**, **show services l2tp session extensive** and **show subscribers accounting-statistics**do not work on LTS (L2TP tunnel switch). PR1596972

User Interface and Configuration

- NETCONF service over SSH with dedicated TCP port (configured with **system services netconf ssh** and the default port is 830) might not work if in-band management is used (that is, connection is established via network interface or loopback interface and so on).PR1517160
- If there is the same configuration stanza across different groups or one of them is in groups, config may not be inherited as expected. PR1529989

VPNs

- In an MVPN environment with the SPT-only option, if the source or receiver is connected directly to the c-rp PE device and the MVPN data packets arrive at the c-rp PE device before its transition to SPT, the MVPN data packets might be dropped. PR1223434
- When ingress PE has duplicate selective tunnel for IPv4 and IPv6, one is wildcard, the other is specific (s, g). If the ingress replication configuration is deleted on the egress PE, sometimes it is observed that the ingress replication entries in ingress PE (DUT) are not properly flushing out for ipv6, but it got flushed out for ipv4, or viceversa. No traffic loss is observed due to this issue. All PIM state and multicast traffic are not impacted due to this issue. PR1475834
- On all Junos OS platforms, if MVPN type 6 and type 7 were intended for local PE router (i.e rt import is local PE address) and later rt import community gets updated (route should be of same NLRI) to remote PE, and further the route itself gets deleted, then it could lead to total c-multicast count to underflow if multiple routes undergoes similar operation. Further MVPN type 6 and type 7 routes may be suppressed and not sent out. PR1567584

SEE ALSO

Resolved Issues

IN THIS SECTION

Learn about the issues fixed 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.

Resolved Issues: 20.1R3

General Routing

- Kernel crash might be seen when micro BFD configuration is applied to a LAG interface. PR1456785
- The MPC2E-NG or MPC3E-NG line card with specific MIC might crash after a high rate of interface flaps. PR1463859
- The following error message is observed after GRES: **[user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, aggIflName:ps1.0 memberIflName:lt-3/0/0.32767]**. PR1466531
- The following line card errors are seen: **HALP-trinity_nh_dynamic_mcast_add_irb_topo:3520 snooping-error: invalid IRB topo/ IRB ifl zero in l2 nh 40495 add IRB**. PR1472222
- Dynamic SR-TE tunnels do not get automatically recreated at the new primary Routing Engine after the Routing Engine switchover. PR1474397

- All VCP interfaces might go down after performing back-to-back VC switchover. PR1480404
- Delay in disabling Packet Forwarding Engine might be seen on MX platforms with MPC7/8/9 and PTX series with PECHIP equipped FPCs inserted. PR1481879
- The SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at the collector. PR1484322
- Traffic impact might be observed in a rare condition when 'monitor traffic interface' is used. PR1498508
- MX204 || Summit || Incorrect log message for PIC1 when changing the configuration from PIC mode to Port Mode. PR1500429
- False positive TSensor errors are reported on vjunos0. PR1508580
- MX150 might go into db mode after software upgrade or downgrade. PR1510892
- Not able to forward traffic to VCP FPC after the MX Virtual Chassis reboots, FPC reboots, or adding VCP link. PR1514583
- On the MX960 routers, the **show interfaces redundancy rlt0** statement shows current status as primary down as FPC is still in the Ready state after RLT failover (restart FPC). PR1518543
- TCP connection going through Packet Forwarding Engine might not be closed at the remote end because there is no TCP FIN segment sent out when the local device is rebooted. PR1517154
- During an upgrade, system displays the following incorrect license warnings when utilizing licensable features even if the license is present on the device: **requires 'idp-sig' license** PR1519672
- Junos OS: Receipt of specific packets could lead to Denial of Service in MQTT Server (CVE-2021-0229) PR1522265
- The bbe-smgd daemon might crash during request repeatedly information of ANCP-triggered out-of-band subscriber sessions. PR1522830
- The BFD session status remains down at the non-anchor FPC even though BFD session is up after the anchor FPC reboots. PR1523537
- No response from the other Routing Engine for the last 2 seconds triggers the following SNMP trap message: Fru Offline. PR1524390
- Packet drops might be seen with all commit events with 1G speed configured interface. PR1524614
- RADIUS Framed-Route sent via RADIUS-Initiated COA message might not be installed into the routing table. PR1524628
- Problem with static VLAN deletion with active subscribers and the FPC might be stuck at the Ready state during restart. PR1525036
- The aggregated Ethernet interface might not come up with LFM configured after reboot. PR1526283
- L2TP subscribers might fail to establish a session on the MX Series device if the CPE is a virtual host PR1527343

- GBUS error on backup Routing Engine is not detected and causes all CBs and FPCs offline after GRES. PR1528319
- Memory leak when querying aggregated Ethernet interface statistics (CVE-2021-0230). PR1528605
- 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
- Commit confirmed doesn't work if the configuration files are encrypted. PR1529499
- Multiple FRUs disconnection alarms might be displayed post the firmware upgrade. PR1529710
- The following error message for port might be observed: **FAILED(-1) read of SFP eeprom**. PR1529939
- The rpd restart might cause prefixes stuck. PR1529965
- It might become unreachable with no console access after performing vmhost reboot post image upgrade. PR1530529
- On MX204 and MX10003 routers, PEM 0 always shows as absent or empty even if PEM 0 is present. PR1531190
- The LACP member link may be down if deleting LFM from it. PR1531235
- On the MX150 routers, configuring the **no-flow-control** command under gigether-options does not work. PR1531983
- Wavelength unlocked alarm is On when using SFP+-10G-T-DWDM-ZR optics. PR1532593
- The rpd might switch to primary role early without syncing up all routes and nexthops if routes churn occurs during GRES switchover. PR1533719
- High rate of host bound traffic on PS interface might lead to eventd high CPU utilization. PR1533721
- **SPD_HA_LOGS_NOT_CONFIGURED: (4882) HA LOGS are not configured 0 0**is seen upon commit when syslog is configured under the service-set. PR1533911
- The dcpfe process might crash and cause FPC to restart due to the traffic burst. PR1534340
- Some routes might get incorrectly programmed in the forwarding table in the kernel, which is no longer present in rpd. PR1534455
- BGP SR-TE IPv6 routes might get hidden after the chassisd restarts. PR1534511
- Multiple vmxt processes might generate core files. PR1534641
- The destination address of ICMP6 neighbor advertisement for VRRP does not comply with RFC 4861. PR1534694
- High rate of ARP or NS packets might be observed between a device that runs Junos OS and host when the device that runs Junos OS receives an ARP or NS packet on an interface in transition. PR1534796
- Snmp mib walk for jnxSubscriber OIDs returns a general error. PR1535754

- All SFBs might go offline due to fabric failure and fabric self-ping probes performing the **disable-pfe** action. PR1535787
- L2PT tunneling fails for VTP upon reboot. PR1536130
- The mixed primary and backup Routing Engine types alarm reported on MX240 with NG-RE. PR1536184
- Specific BGP VPNv6 flowspec message causes routing protocol daemon (rpd) process to crash with a core (CVE-2021-0236). PR1537085
- The chassisd memory leak might cause traffic loss. PR1537194
- Deactivating or activating PTP or synchronized Ethernet in the upstream router causes the 100GbE links on the LC2103 to flap. PR1538122
- Traffic drop might be seen while executing the **request system reboot** command. PR1538252
- The login of subscriber might fail. PR1538971
- The rpd memory leak might be observed on the backup Routing Engine due to the flapping of the link. PR1539601
- Some telemetry paths do not work.. PR1539603
- Any layer 2 traffic with ether type 0x88b7 might get punted to Routing Engine and dropped on MX platforms. PR1539807
- The MPC may show 'Absent' after GRES. PR1540068
- FPC might not be recognized after the power cycle (hard reboot). PR1540107
- BPDUs are not sent out when the interface on anchor FPC is down. PR1540380
- The mspmand process leaks memory in relation to the MX Series telemetry reporting the following error message: **RLIMIT_DATA exceed**. PR1540538
- Subscriber might not come up on some dynamic VLAN ranges in a subscriber management environment. PR1541796
- The KRT queue might get stuck after the Routing Engine switchovers. PR1542280
- Port mirroring with the maximum-packet-length configuration does not work over the GRE interface. PR1542500
- The JNH memory leak could be observed on MPCs or MICs. PR1542882
- Any unsupported timing related configuration may cause the FPC to crash on MX10008 and MX10016. PR1543742
- Backup Routing Engine boots up as master after upgrade or reboot when chassis redundancy failover configuration statements are enabled. PR1543879
- Phy-Sync state toggles when introducing phase-offset in the syncE line in the hybrid scenario, thereby affecting the downstream slaves. PR1543881

- Sessions creation rate is set to minimal rate after IDS and CPU throttling in place during DDOS attack. PR1544489
- ANCPD core when hitting maximum-discovery-table-entries limit (CVE-2021-0224). PR1544746
- The mspmand process might generate the core file on activating or deactivating the interface. PR1544794
- The kmd process might crash when the interface flaps. PR1544800
- Traffic loss might be observed when the Switch Fabric Board 3 and MPC8E 3D combination is used in the MX2010 and MX2020 routers. PR1544953
- The chip on FPC line card might crash when the system reboots. PR1545455
- Continuous rpd errors might be seen and new routes fails to be programmed by the rpd process. PR1545463
- **jnxDomAlarmSet** and **jnxDomAlarmClear** trap are getting generated at 15 minutes intervals after a link on the transceivers support DOM becomes up or down. PR1545514
- FPC(s) may not boot-up on MX960 in a certain condition PR1545838
- The performance of Packet Forwarding Engine process on MX204 might be degraded. PR1545989
- The receipt of specific DHCPv6 packet may cause jdhcpd to crash and restart (CVE-2021-0241). PR1546166
- Unexpected log messages appears related to the Neighbor Solicitation (NS) messages with multicast as source address. PR1546501
- The traffic of MPLS-IPv4 FEC might not be forwarded properly in JFlow or inline-JFlow scenario. PR1546600
- Backup Routing Engine vmcore might be seen due to the absence of the next-hop acknowledgement infra. PR1547164
- The PTP protocol might get stuck at Initializing state on MX platforms. PR1547423
- 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 the inline NAT in the USF mode. PR1547647
- Traffic for some v4ov6 entries is dropped. PR1547681
- The **SENSOR APP DWORD** leak is observed during the period of churn for routes bound to the sensor group. PR1547698
- Multicast traffic drop might be seen after ISSU. PR1548196
- The MS-MPC/MS-MIC located at VC-B might not work properly in an MX Series Virtual Chassis. PR1548340
- FPC crash may occur after flapping the multicast traffic. PR1548972
- The KRT queue might get stuck after Routing Engine switchover if NSR is enabled. PR1549345

- The rpd process might crash if performing multiple GRES. PR1549884
- PKI CMPv2 client certificate enrolment does not work on SRX when using root-CA. PR1549954
- The LLDP adjacency could not be established for fxp interface. PR1550131
- MPC crash might happen on MX platforms. PR1550575
- Two Routing Engine's might lose communication if they have different Junos OS versions on MX10003. PR1550594
- The adapted sample rate might get 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 crash might be seen when the BGP service route is resolved over the color-only SR-TE policy. PR1550736
- The l2alm processes high CPU utilization might be observed in the EVPN-VXLAN environment. PR1551025
- The Routing Engines may crash during swapping the child interfaces of two aggregated Ethernet bundles. PR1551184
- The PPPoE subscribers might fail to login. PR1551207
- The rpd crash might be seen when **multipath-resolve** and **preserve-next-hop-hierarchy**are both configured. PR1551308
- The IRB interface might not work after chassisd and l2ald reboot in EVPN scenario. PR1551631
- The **LCM Peer Absent** message might be seen. PR1551760
- The softwire might not be established when connecting to a different AFTR. PR1552431
- The statement 'action-shutdown' of storm control does not work for ARP broadcast packets. PR1552815
- The fabric errors are observed and the FPC processes might get offline with the SCBE3, MPC3E-NG, or MPC3E and MPC7 or MPC10 line card in the increased-bandwidth fabric mode. PR1553641
- ISSU upgrade from pre 19.1 to 19.1 onwards may cause few interfaces to go down. PR1554099
- The link on the Linux based LC is not brought down immediately after the FPC process(ukern/indus.elf) crashes or the process is killed. PR1554430
- During ISSU, BNG losses subscriber sessions without sending the Session Stop message but stay in authd. PR1554539
- The device takes 3/10 mins to bring up the 100/1000 subscribers. PR1555216
- The VGA might be down when configuring the IRB interface with multi VGA addresses. PR1555338
- The LCMD process might consume memory until all of the free memory available to VMHOST gets exhausted. PR1555386
- The subscriber's RADIUS interim accounting statistics update might not work in some scenario. PR1555492

- MACsec session may remain down after CA applied or re-applied to different interfaces PR1555736
- Fabric self ping failure might be reported from MPC10 when MPC CPU is busy. PR1555802
- The following message is not generated on the MPC11E line card due to no power: Chassisd SNMP trap Fru Offline. PR1556090
- Junos OS: FPC may crash upon receipt of specific MPLS packet affecting Trio-based MPCs (CVE-2021-0288). PR1556576
- The framed route installed for a demux Interface has no MAC address. PR1556980
- The framed-routes are stuck in "KRT queued (pending) add" state when "routing-service enable" is configured under dynamic-profile. PR1557230
- Multiple FPCs crash might be seen when performing GRES or FPC reboot repeatedly in subscriber scenario. PR1557294
- ISSU are aborted and the chassisd process generates core file on the backup Routing Engine during the Junos OS upgrade to version Junos OS Release 20.2R2-S1. PR1557413
- Packets corruption on 100G or 40G interface are configured with protocol PTP. PR1557758
- On the MX150 routers, the following continuous license error is observed: **[licinfra_set_usage_nextgen_async:1733] Invalid input parameters**. PR1559361
- The subscriber management infrastructure daemon (smid) process might be stuck at 100%. PR1559402
- The PTP master line card servo might stuck in Freerun state. PR1560074
- The "jnxDomAlarmSet" and "jnxDomAlarmClear" trap will be generated for a copper port. PR1560149
- The **request system software validate** command might corrupt installation of the junos-openconfig package. PR1560234
- The VXLAN queue DDos violation and RARP packets flood might happen if receiving the RARP packets more than the supported DDoS bandwidth. PR1560243
- The tunable optics SFP+-10G-T-DWDM-ZR does not work. PR1561181
- The l2cpd process might generate a core file on reboot. PR1561235
- Continuous bbe-smgd cores are generated after restarting the smgd. PR1561855
- Routing Engine switchover on-disk-failure does not work as expected when GRES is disabled. PR1563505
- Failure in validating packages in seen when MS-MPC PIC is booted up in FIPS mode. PR1562202
- The following error message might be seen after ISSU: Turbotx process not running. PR1564418
- MX platforms with MX-SCBE3 may reboot continuously. PR1564539
- Junos OS: Upon receipt of specific packets BFD sessions might flap due to DDoS policer implementation in Packet Forwarding Engine (CVE-2021-0280). PR1564807
- Commit error observed when tunnel-service is configured on a PIC without explicit bandwidth. PR1565034

- On the MX2010 or MX2020 routers, the following error message might be observed after switchover with **GRES/NSR: CHASSISD_IPC_FLUSH_ERROR**. PR1565223
- Unable to bring up more than one client on one VLAN at the same time. PR1565249
- PPPoE service-name-tables does not correctly count active sessions that matches agent-specifier aci/ari used for delay. PR1565258
- The KRT log file might continue to grow after removing the KRT log configuration. PR1565425
- The mspmand crash might be seen on the PIC of MS-MPC/MS-MIC. PR1566325
- The chassisd crash might be seen on MX platforms. PR1567479
- TLB composite NH is installed incorrectly in other routing-instances. PR1567568
- Need to allow the tunnel interface as the peer-address for ALQ. PR1567735
- On the MX204 routers, FPC might display high CPU utilization because of the JGCI background thread that 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
- BFD flaps seen between leaf and core during spine reboot causing other protocols flap. PR1568615
- SPC3 card interfaces are not created. PR1568694
- Traffic loss might be observed when SCU accounting is configured and logical-systems is enabled. PR1569047
- The agent sensor __default_fabric_sensor__ are partly applied to some FPCs, which causes zero payload issue AGENTD received empty payload for pfe sensor __default_fabric_sensor__. PR1569167
- Junos OS and Junos OS Evolved: LLDP Out-of-Bounds Read vulnerability in l2cpd (CVE-2021-0277). PR1569312
- Wi-Fi mPIM on SRX Series devices is reaching out to NTP and DNS servers. PR1569680
- The MPLS traffic passed through the back-to-back PE topology might match the wrong CoS queue. PR1569715
- The mspmand process might crash if the packet flow-control issue occurs on MS-MPC/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 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
- The TFEB/FPC may fail to be online after rebooting the system or the FPC if interface-set is configured for CoS PR1572348

- On the MX960 routers, the 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
- 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 the MPC3-NG/MPC3E/SRX5K-IOC2 line cards are installed along with the MPC7/MPC10/SRX5K-IOC04 and SCBE3/SCB4 line cards operating in an increased-bandwidth fabric mode. PR1573360
- QSFP 4x10G interface might not come up after FPC reboot. PR1574279
- Slow FPC heap memory leak might be triggered by flapping the subscribers terminated over multiple pseudowires. PR1574383
- PTP might be stuck in Phase acquiring state after ISSU upgrade. PR1575055
- On the MX150 routers, the interface might take a long time to power down while rebooting, powering-off, halting, or upgrading. PR1575328
- On the 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/SPC3 might reset on receiving the Subscriber Traffic. PR1576946
- Commit failure-error: Modified IFD "ae0" is in use by targeted Broadband Edge subscriber, commit denied - mtu config changed (1522), (1514). PR1577007
- Traffic loss might be seen when subscriber service over aggregated Ethernet bundle interface(s). PR1577289
- When line card is booted on RE1 being Master, Nextgen stats failed to fetch the value of backup mac address correctly. PR1577611
- MX-VC: gRPC based /components/ sensor output is missing lot of data. PR1580120
- More than one subscriber on same vlan fails to apply same FWF Template. PR1580826
- kern.ipc.maxpipekva exceeded; see tuning error. CLI does not display output. PR1581192
- hitting with vmcore.0 at 0xffffffff80443eef in kern_reboot PR1581260
- The rpd process might crash on the new master after performing graceful switchover. PR1581878
- Changing bandwidth statement does not take affect for SNMP ifHigSpeed oid until a PSX interface is disable or enabled. PR1582060
- The voice VLAN might not get assigned to the access interface. PR1582115
- bbe-smgd process on both routing engines may crash due to a rare timing issue after logout of subscribers over pseudowire. PR1582356
- SNMP SysObjectID.0 is empty with enabled unified-services. PR1583534

- TCP connection to syslog server might fail to be established after adding **tcp-log** configuration for an existing service-set. PR1583979
- ATT VAN: jsd process hogging CPU. PR1584357
- The rpd 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 in **show bridge statistics instance**. PR1584874
- After changing configuration, **show bridge statistics** shows extreme larger value. PR1584876
- The bbe-statsd memory leak might be observed on backup Routing Engine during subscribers login or logout. PR1589081
- Traffic loss might be observed for interface configured in subnet 137.63.0.0/16. PR1590040

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. PR1556103
- FPC crash might be observed after the **show class-of-service** command. PR1568661

EVPN

- With dynamic list next hop configured, a forwarding problem occurs after performing graceful switchover. PR1513759
- no-arp-suppression is required for MAC learning across the EVPN domain on the static VTEP. PR1517591
- The local PE does not remove VNI flood information even though it does not receive VXLAN message from remote PE PR1520688
- The kernel crash might happen in EVPN-VXLAN scenario PR1524485
- The BUM (Broadcast, Unknown Unicast, and Multicast) 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 Junos OS Release19.4R1 S2. PR1534021
- All the ARP reply packets toward some address are flooded across the entire fabric. PR1535515
- Rpd memory leak might occur when the EVPN configuration is changed. PR1540788
- The rpd process 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 in overlayd service (CVE-2021-0254). PR1548415
- Missing ARP entry might happen in certain conditions under EVPN Active-Standy scenario. PR1550577

- The rpd might crash under EVPN-VPWS environment. PR1562160
- Policy with mac-filter-list might not work if change unrelated to that policy is committed in EVPN scenario. PR1567623
- The multicast traffic loss might be seen in EVPN-VXLAN scenario with CRB multicast snooping PR1570883
- The mustd.core process generates core file during upgrading or while committing a configuration. PR1577548
- RPD might crash in high scaled EVPN-VXLAN scenario. PR1581674
- The BUM traffic might lose after triggering NSR in EVPN-MPLS or EVPN-ETREE scenario. PR1586402

Forwarding and Sampling

- The l2ald process might crash when a device configuration flaps frequently. PR1529706
- MAC learning issue might happen when EVPN-VXLAN is enabled. PR1546631
- All traffic are dropped on the aggregated Ethernet interface bundle without the VLAN configuration if the bandwidth-percent policer is configured. PR1547184
- The l2ald process might crash due to next-hop issue in the EVPN-MPLS. PR1548124
- The dfwd process might crash when implementing non-contiguous firewall filter. PR1555724
- The configuration archive transfer-on-commit fails when running Junos OS Release 18.2R3-S6.5. PR1563641
- In the VXLAN scenario, the locally originated packets have UDP source port 0. PR1571970
- The pfed memory leak might be observed. PR1573285

General Routing

- DHCPv6 subscribers are not synchronized on redundancy DHCP relay agents with ALQ enabled. PR1538644
- 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
- The DHCP ALQ is not working as expected. PR1578543

Infrastructure

- Invalid statistics value might be observed when multiple mib2d/cosd requests for the same IFD arrive within 1 second. PR1541579

Interfaces and Chassis

- The vrrpd might crash when dual VLAN on VRRP interfaces is configured. PR1512658
- Junos OS: ethtraceroute Local Privilege Escalation vulnerabilities in SUID binaries (CVE-2021-0255). PR1529209
- The configuration might not be applied after deleting all existing logical interfaces and adding a new logical interface for an IFD in a single commit. PR1534787
- After VRRP failover, the VRRP backup router keeps receiving traffic for about 2 minutes. PR1546635
- The following commit error is observed while trying to delete unit 1 logical system interfaces: ae2.1: Only unit 0 is valid for this encapsulation. PR1547853
- An irb interface which has large unit value over 32767 cannot be active group for inheriting VRRP. PR1550993
- 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
- Junos device might send VRRP advertisement packets in the VRRP Init or Idle state before startup-silent-period timer expiry on the VRRP primary device with NSR disabled after GRES. PR1558560
- The **input errors counter** command on the **monitor interface** command does not work. PR1561065
- MAC address entry issue might be observed after the MC-LAG interface. PR1562535
- if-media-type missing from interface XML output on MX platforms. PR1574035
- MC-AE interfaces may go down if same VRRP group-id is configured on multiple IRB units. PR1575779
- Unable to configure pseudowire interface on an MX10003 in virtual chassis mode. PR1587499

Layer 2 Ethernet Services

- The jdhcpd crash and ALQ sync issue might be seen after rebooting an ALQ device. PR1552855
- DHCP packet drop might be seen when the DHCP relay is configured on a leaf device. PR1554992
- In a DHCP relay configuration with active lease query, some subscriber's active on master might get logged out. PR1559269
- Receipt of malformed DHCPv6 packets causes jdhcpd to crash and restart. (CVE-2021-0240) PR1564434
- DHCPv6 Option 18 and Option 37 might not be created in DHCP dual stack scenario. PR1564778
- The jnxJdhcpLocalServerMacAddress (.1.3.6.1.4.1.2636.3.61.61.1.4.3) returns incorrect format of the MAC address. PR1565540
- 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
- DHCP relay drops packets during the renewal DHCP process. PR1576417
- The jdhcpd might crash if "relay-source lo0" enabled in DHCP relay. PR1580724
- The jdhcpd process might not respond to any Discover message when it is in "clients waiting to be restored" state. PR1592552

MPLS

- The rpd scheduler might slip after the link flaps. PR1516657
- LDP Entropy Label Capability might be always enabled for routes from SR-to-LDP stitching in the scenario of interoperability of Segment Routing with LDP. PR1545404
- If link-protection is enabled for an externally provisioned LSP, any commit for the first time after provisioning causes a break (MBB) even if the configuration is not related to the LSP. PR1546824
- A new LSP might not be up even if bypass LSP is up and setup-protection is configured. PR1555774
- Incorrect EXP bit change might be seen in certain conditions under MPLS scenario. PR1555797
- MPLS-LIB memory leak might be seen in SR scenario. PR1556495
- LDP routes might be stuck when BGP LU session is down. PR1562884
- Traffic sent over an LSP may be dropped if two consecutive PLRs along the LSP perform local repair and bypass protecting the second PLR fails PR1566101
- The rpd process on the transit node might crash when MPLS traceroute on the ingress node is performed. PR1573517

Multicast

- FPC might crash in a multicast scenario. PR1569957

Network Management and Monitoring

- The SNMP trap sent from snmpd might be dropped silently. PR1530830
- Commit error occurs while deleting the routing instance when the SNMP trap-group also have the same routing instance referred. PR1555563
- The mib2d process crashes and generates a core dump on backup Routing Engine. PR1557384

Platform and Infrastructure

- PE-CE OAM CFM might have issues in the aggregated Ethernet interface. PR1501656
- Interwork failure between Junos OS Evolved as RPM client and TVP platforms as RPM server (and vice versa). PR1508127
- MX Series: DDoS LACP violation upon receipt of specific layer 2 frames in EVPN-VXLAN deployment (CVE-2021-0228). PR1512033

- The state of the flow detection configuration might not be displayed properly if DDoS-SCFD is configured globally. PR1519887
- Flow programming issue for lt- interface in the Packet Forwarding Engine level is observed. PR1525188
- The following error message is observed when alarms after interface reset: **7836 ifl 567 chan_index 8 NOENT & jnh_ifl_topo_handler_pfe(13015): ifl=567 err=1 updating channel table nexthop**. PR1525824
- The rpd process might crash in a rare timing condition with logical system. PR1527242
- Packet Forwarding Engine errors 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 the 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 on VPLS in the PS interface. PR1536043
- COS queue egress interface forwarding-class might not work as expected. PR1538286
- Packet loss might be observed when the RFC2544 egress reflector session is configured on the non-zero Packet Forwarding Ethernet interface. PR1538417
- The following major error message might cause the Packet Forwarding Engine(s) to disable: **XQ_CMERROR_SCHED_L3_PERR_ERR**. PR1538960
- Subscribers over an "interface-set" may not be able to log on. PR1539260
- The VPLS traffic loss might be observed when member0-FPC0 is not alive in MX VC scenario. PR1539562
- Trio-based 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
- The RP expired timer on the backup Routing Engine is not the same as the primary Routing Engine if the aging-timer is configured. PR1544398
- Junos OS: EX4300: FPC crash upon receipt of specific frames on an interface without L2PT or dot1x configured (CVE-2021-0242). PR1545530
- In rare occurrence Routing Engine kernel might crash while handling TCP sessions if GRES/NSR are enabled. PR1546615
- An internal timer on the backup Routing Engine might cause an ARP storm upon GRES switchover on the new primary (old backup) Routing Engine. PR1547583
- The IIF-LIST APP DWORD leak is observed during the period of churn for the NGMVPN-MoFRR routes with sender-based-rpf enabled. PR1548806
- The kernel might crash if GRES is performed on either new iteration or after swapping the Routing Engine and restoring the HA configuration. PR1549656
- FPC might crash due to the errors triggered by microcode performance optimization operation in GUMEM. PR1550933

- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. PR1552603
- ARP resolution might fail if ARP packets are received over multicast based VxLAN access network from CE. PR1553917
- FPCs may go to "ISSU error" state post performing enhanced ISSU. PR1553961
- Traffic is not forwarded over IRB to a Layer 2 circuit on the lt interfaces. PR1554908
- The IPv4 EXP rewrite might not work properly when inet6-vpn is enabled. PR1559018
- The BUM (Broadcast, Unknown Unicast, and Multicast) 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
- The DHCPv4 request packets might be wrongly dropped when DDoS attack occurs. PR1562474
- The mtr process might hog CPU when the **traceroute monitor** command is paused. PR1563298
- The **enforce-strict-scale-limit-license** configuration enforces subscriber license incorrectly in the ESSM subscriber scenario. PR1563975
- "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 Packet Forwarding Engine.PR1567515
- pfe err-jnh_physmem_add_resvd_to_cntr(18014): PFE 0 jnh_app 0x08020860, add ox00080000 from 0x00b00000-0x00b80000 to baMask 0x1. PR1570631
- 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 exhausted the FPC might reboot. PR1571439
- Scale-subscriber license might be not updated properly on the backup RE which leads to "License grace period for feature scale-subscriber(44) is about to expire" alarm after GRES. PR1573289
- cassxr_err_addr(8593): Uninitialized Read Error @ 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 PFE after applying sampling and "flex-flow-sizing" to the Trio based line-cards. PR1575994
- VRRP device originally taking slave role might cause destination IP unreachable after VRRP mastership switch-over. PR1584115
- Authorization issue might be observed if the login class is configured on NFX platforms. PR1586377
- FPC might crash in a scaled-firewall configuration. PR1586817

Routing Policy and Firewall Filters

- The rpd might crash continuously if "then install-nexthop" policy is applied to a route pointing to list nexthop. PR1488818
- 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
- Global variable *policy_db_type* is not set to the correct value on failure. PR1561931
- Generated route goes to the **Hidden** state when the **protect core** command is enabled. PR1562867
- The rpd might crash due to 'source-address-filter-list' enabled within the policy. PR1565891
- **bbe-smgd - dymanic-profile NACK** due to config error reading address mask prefix-length in policy-options/policy-statement. PR1583535

Routing Protocols

- The BGP RPKI ROA withdrawal might lead to an unexpected BGP route flap. PR1483097
- Virtual Chassis or Virtual Chassis fan, the following error message is observed while copying the image to the Virtual Chassis fan member and trying to downgrade the image: rcp for member 14, failed. PR1486632
- The rpd might crash with BGP RPKI enabled in a race condition. PR1487486
- High CPU utilization by BGP I/O thread on master RE might be seen if NSR is enabled on a large-scale BGP setup. PR1488984
- Traffic might be silently discarded when the BGP route gets deleted, which is part of multipath. PR1514966
- **mld snooping membership** command is not accepting group with vlan options together (Same works with IGMP snooping membership command). PR1516650
- BGP session might not advertise routes to peer in a VRF instance. PR1522049
- The BGP session with VRRP virtual address might not come up after a flap. PR1523075
- 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
- The rpd process generates the core file at **gp_rtarget_tsi_update,bgp_rtarget_flash_rt,bgp_rtarget_flash**. PR1541768
- Continuous rpd crash might be observed if a static group is added to protocol PIM. PR1542573
- ISIS route convergence from L1 to L2 might take more than 10 minutes. PR1542932
- 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

- Unexpected packet loss may happen due to inet-vpn routes not valid in vrf.inet.0 and bgp.l3vpn.0 routing tables. PR1543717
- If output-queue-priority expedited update-tokens is configured, rpd might crash might upon BGP flapping. PR1545837
- With BGP rib-sharding enabled, the RPD memory exhaustion might be observed. PR1546347
- VPLS multicast traffic loss might be observed after RE switchover. PR1546397
- BGP session might be down due to BGP-LS TLV received out of order. PR1546416
- The rpd memory leak might be seen in the BGP scenario. PR1547273
- BGP convergence delay may occur in a scale BGP setup. PR1548517
- The BGP session neighbor shutdown configuration does not effect the non-established peer. PR1554569
- The changes do not get effective when the values are set under the static default hierarchy. PR1555187
- Sending multicast traffic to downstream receiver on Trio based Virtual Chassis platforms might fail. PR1555518
- Junos OS and Junos OS Evolved: RPD could crash in SR-ISIS/MPLS environment due to a flap of an ISIS link in the network (CVE-2021-0287). PR1555627
- The rpd crash might be seen when removing/adding BGP configuration. PR1556062
- Junos OS and Junos OS Evolved: Specific packets can trigger rpd crash when BGP Origin Validation is configured with RPKI (CVE-2021-0281). PR1556207
- The rpd core might occur when BGP origin validation trace is enabled with scaled routes. PR1556210
- Six PE device prefixes might not be removed from RIB upon reception of withdrawal from a BGP neighbor when the RIB sharding is enabled. PR1556271
- Prefix learning issue might be seen if ISIS with "topologies ipv6-unicast" is enabled. PR1557726
- BGP LU session flap might be seen with AIGP used scenario. PR1558102
- The ISO routes are not leaked in default (master) instance after switchover or reconfiguration. PR1558532
- Traffic loss might occur for stitched traffic from SR towards LDP if no-eligible-backup is configured. PR1558565
- The rpd process might crash when applying the BGP route policy change. PR1560037
- All the Layer 3 VPN route resets when a VRF is added or removed. 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 the interface flaps. PR1561414
- The ppmd memory leak may cause traffic loss. PR1561850

- The rpd process might crash if there are more routes changed during the commit-sync processing window PR1565814
- There might be traffic loss when GRE interface flaps on QFX platforms. PR1566428
- The rpd process might crash in BGP L2VPN scenario due to memory corruption. PR1567026
- The rpd process might crash when there is BGP session re-establishing or flapping. PR1567182
- The rpd memory leak may be observed during CLI/ephemeral commits in OSPFv2 scenario. PR1568157
- Traffic might be lost during mirror data transmit from the primary ppmd or bfdd. PR1570228
- SNMP MIB ospfv3NbrState is returning drifted value. PR1571473
- 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 decapsulation interface is not created as expected. PR1577461
- The rpd may crash when two or more routing instances are deleted in one shot. PR1578740
- The rpd might crash in BGP and MPLS scenario. PR1581794
- The dcpfe process might crash when any interface flaps. PR1579736
- Authentication may fail if the password contains special characters. PR1580003
- With IGMP snooping implemented, there is unexpected jitter issue that could cause traffic loss. PR1583207
- The rpd crash might be seen after committing with static group 224.0.0.0 configured. PR1586631
- Wrong BGP next-hop advertisement in L3VPN scenario. PR1587879
- The routing process may crash due to memory corruption while processing BGP multipath route. PR1594626

Services Applications

- The kmd might be stuck when NAT is configured on the device between IPsec gateways. PR1494473
- Junos OS: MX Series: Executing CLI command repetitively may cause the system to run out of disk space (CVE-2021-0238). PR1537772
- L2TP tunnel might fail after L2TP silent failover is triggered on the L2TP LNS node. PR1541122
- The following error message is observed: **SPD_CONN_OPEN_FAILURE: spd_pre_fetch_query: unable to open connection to si-1/0/0**. PR1550035
- Protocol flapping might happen in IPsec SA enabled scenario. PR1550407
- The CoA with LI-on/LI-off message might be dropped during CoA process. PR1554618

- Memory leak might be observed in tunnel flapping scenario. PR1567291
- IWF AVP value may not be reflected properly on LTS. PR1581096

Subscriber Access Management

- The BNG authd process memory leak might happen after the subscriber logout. PR1530820
- Only the last RADIUS class attribute is considered when sending more than one attribute at one time. PR1536129
- PCRF might not work properly if subscriber logging in is dual stack subscriber. PR1545307
- The authd may crash after performing ISSU in MX BNG scenario. PR1570096
- CoA request may not be processed correctly from time to time. PR1571501

User Interface and Configuration

- Any change in the nested groups might not be detected on commit and does not take effect. PR1484801
- Commit might fail after the Routing Engine switchovers PR1531415
- The verbose command unexpectedly becomes hidden after Junos OS Release 16.1 for **set system export-format json**. PR1547693
- The chassisd core dump might be observed if PIC number 2 or 3 is used on MX204. PR1555685
- Apply-paths might cause validation failures during JUNOS upgrade. PR1577626

VPNs

- The l2circuit local-switching end interface might get stuck in XX (Unknown) state upon vlan-id-list configuration change. PR1528809
- MVPN multicast route entry might not be properly updated with the actual downstream interfaces list. PR1546739
- Type7 messages may not be sent from egress PE resulting in Type 3/5 messages not created for some S, Gs in source PEs. PR1567584
- The rpd might crash in the NG-MVPN scenario on all Junos and Junos EVO platforms. PR1579963
- The ddos-protection reason "packets failed the multicast RPF check" may be seen in NG-MVPN scenario with GRE transport. PR1591228

Resolved Issues: 20.1R2

Application Layer Gateways (ALGs)

- FTPS traffic might be dropped on MX Series platforms if FTP ALG is used. PR1483834
- The srxpfe and mspmand processes might crash if FTPS is enabled in a specific scenario. PR1510678

Class of Service (CoS)

- SNMP query for jnxCos objects does not work. PR1475960
- The following error message is observed: **GENCFG write failed (op, minor_type) = (delete, Scheduler map definition) for tbl id 2 ifl 0 TABLE Reason: No such file or directory**. PR1476531
- The MX Series routers with MPC1 Q and MPC2 Q line cards might report memory errors. PR1500250

EVPN

- When a dynamic list next hop is referenced by more than one route, it might result in an early deletion of the next hop from the kernel, thereby assigning the next-hop index as 0. PR1477140
- IRB interface might get stuck in down state in an EVPN multihoming scenario. PR1479681
- Deleting a Layer 2 logical interface generates an error if the interface is not deleted first from EVPN. PR1482774
- The ESI of IRB interface does not get updated after autonomous-system number change if the interface is down. PR1482790
- Due to timing condition, dead next hops in the flood group of the EVPN-MPLS are seen after remote PE devices bounce. PR1484296
- The ARP entry gets deleted from the kernel after adding and deleting the virtual gateway address. PR1485377
- The rpd process might crash due to a slow memory leak. PR1490269
- The rpd process might generate a core file when the Routing Engine switches over after disabling the BGP protocol globally. PR1490953
- VXLAN bridge domain might lose the VTEP logical interface after restarting chassisd. PR1495098
- The l2ald memory leakage might be observed in any EVPN scenario. PR1498023
- Packets might not be sent out of the IRB interface if there is no Layer 2 interface in the associated bridge domains. PR1498534
- In an EVPN-VXLAN scenario, the l2ald process might crash in a rare condition. PR1501117
- The VXLAN function might be broken due to a timing issue. PR1502357
- The MAC address of the LT interface might not be installed in the EVPN database. PR1503657
- Configuring the **proxy-macip-advertisement** statement for EVPN-MPLS leads to functionality breakage. PR1506343
- Unable to create a new VTEP interface. PR1520078
- ARP table might not be updated after performing VMotion or a network loop. PR1521526
- EVPN: routing table stuck in deleted state in kernel. PR1521668
- The BUM traffic might get dropped in the EVPN-VXLAN setup. PR1525888

- The rpd process might crash when **auto-service-id** is configured in the EVPN-VPWS scenario. PR1530991
- All the ARP reply packets toward to some address are flooded across the entire fabric. PR1535515

Forwarding and Sampling

- In some rare scenarios upon FPC or PIC reboot, the Packet Forwarding Engine daemon database might not get updated with the correct location_id for some physical interfaces. Then a problem with statistics on some interfaces of a router might be observed. PR1458143
- DHCP relay might not work normally under an EVPN with VXLAN environment. PR1487385
- The DHCP subscribers might get stuck in Terminated state for around 5 minutes after disabling cascade ports. PR1505409
- The pfed process might crash while running the **show pfe fpc x** command. PR1509114
- UTC timestamp is used in the flat-file-accounting files when a profile is configured. PR1509467
- Traffic might be dropped without exceeding the configured bandwidth under policer. PR1511041
- The srrd process might crash in a high route churns scenario or if the process flaps. PR1517646

General Routing

- In some MX Series platforms, the following random syslog messages are observed for FPCs: **fpcx ppe_img_ucode_redistribute Failed to evict needed instr to GUMEM - xxx left**. PR1298161
- The **show security group-vpn member IPsec security-associations detail | display xml** command is not in the expected format. PR1349963
- The **max-drop-flows** statement is not available. PR1375466
- On the MX2000 router, the error message **Failed to get xfchip** might be observed if the MPC7 line card is offline when Routing Engine switchover occurs. PR1388076
- After the JNP10K-LC2101 line card is powered on, a voltage of 1345–1348 mV is read for about 20 seconds, which gets stabilized to 1493 mV. During this period, the **FPC x Voltage Tolerance Exceeded** major alarm is raised. PR1415671
- In some scenarios with PTP hybrid mode, continuous resetting of the Playback Engine log message occurs. PR1420335
- FPC might crash after GRES when you commit the changes in firewall filter with the **next term** statements in the subscriber scenario. PR1421541
- PTP might not work on the MX104 router with any two-port license installed on the 10GbE interface and if phy-timestamping is enabled in PTP. PR1421811
- The RPD scheduler slips might be seen upon executing the **show route resolution extensive 0.0.0.0/0 | no-more** command if the number of routes in the system is large (several millions). PR1425515
- PTP and show warning are disabled when hyper mode is configured. PR1429527
- MPC9E line card does not go offline due to unreachable destinations in phase 3 stage. PR1443803

- FEC statistics are not getting reset after changing the FEC mode. PR1449088
- The **Mixed Master and Backup RE types** alarm is observed when MX2008 with RE-MX2008-X8-128G detects backup Routing Engine as RE-MX2008-X8-64G. PR1450424
- When an M-VLAN interface (OIF map) is changed, the existing multicast subscribers with membership reports in place experience loss of multicast traffic till traffic is forwarded to the new OIF map. For example, a new M-VLAN interface. PR1452644
- Interfaces shut down by the **disable-pfe** action might not come up when you use the MIC offline or online command. PR1453433
- The FPC or the Packet Forwarding Engine might crash with the ATM MIC installed in the FPC. PR1453893
- When the scale configurations are applied, the chassisd CLI command might delay response or might time out for 10 minutes. PR1454638
- Application and removal of 1-Gbps speed results in the channel being down. PR1456105
- LSP statistics are not getting reset after routing restart. PR1458107
- Multiple leaf devices and prefixes are missing when an LLDP neighbor is added after streaming is started at the global level. PR1460347
- In the MVPN instance, traffic drops on multicast receivers within the range of 0.1 to 0.9 percent. PR1460471
- On the MX960 router, the following error message might be observed: **SCHED L4NP[0] Parity errors**. PR1464297
- The bbe-smgd process generates core files on the backup Routing Engine. PR1466118
- On the MPC11E line card, the DOM MIB alarm for the channelized 10-Gigabit Ethernet interface does not show any alarm for LF/RF. PR1467446
- In Junos OS Release 16.2R1 and later, if commit is executed after commit check, the daemons (for example, dhcpd and sampled) might not get started even after the related configuration is successfully committed. PR1468119
- On the MX150 routers, the request system halt and request system power-off commands do not work as expected. PR1468921
- The GRE tunnel might go down in a scenario with IPv4 and IPv6 IPsec service configured. PR1470667
- The following syslog message are observed: **fpcX user.notice logrotate: ALERT exited abnormally with [1]**. PR1471006
- On MX104 routers, the clksyncd crash might be seen when PTP over an aggregated Ethernet is configured. PR1471466
- pkid process might crash at bn_i2c (pval=0x1d, cont=0x0, putype=0xffffcce8, it=0xc8c848b8 <BIGNUM_it>) at ../../../../../../../src/crypto/openssl/crypto/asn1/x_bignum.c:127. PR1471878

- When both MSTP and ERP are enabled on the same interface, ERP might not come up properly. PR1473610
- Drops counter does not increment for the aggregated Ethernet interface even after the member link shows the drops. PR1473665
- On MX150 platforms, core files are not seen under show system core-dumps. PR1474118
- A newly added LAG member interface might forward traffic even though its micro-BFD session is down. PR1474300
- Traffic loss might be seen as backup Routing Engine takes around 20 seconds to acquire the primary role. PR1475871
- Syslog reports simultaneous zone change reporting for all zones green, yellow, orange, red for one or more service PICs. PR1475948
- On MX2020, MX2010, and MX960 platforms, traffic drop might be observed while performing a unified ISSU. PR1476505
- In vMX instances, after every commit, the following error message is observed: **chassisd[7836]: %DAEMON-3-CHASSISD_IOCTL_FAILURE: acb_get_fpga_rev: unable to get FPGA revision for Control Board (Inappropriate ioctl for device)**. PR1477941
- The ukern-platformd process might crash on the MX2000 router with the MPC11 line card. PR1478243
- The FPC with **vpn-localization vpn-core-facing-only** configured might be stuck in ready state. PR1478523
- On the MPC7E, MPC8E, and MPC9E line cards, hardware sensor information is logged on the syslog and /var/log/messages for every 30 minutes. PR1478816
- All PPPoE subscribers might not log in after FPC restart. PR1479099
- Multiple SQLite vulnerabilities resolved. PR1480208
- The 100GbE might randomly fail to come up after maintenance operations. PR1481054
- Memory utilization enhancement is required. PR1481151
- Issue with binding non-default routing instance to existing soft-gre group. PR1481278
- After unified ISSU on the primary and the backup Routing Engines, the **ISSU enhanced-mode:Performing action get-state for error /fpc/5/pfe/0/cm/0/PCIe_Error/0/PCIE_CMERROR_UNCORRECTABLE (0x190001)** error message is generated. PR1481859
- Fabric healing logic incorrectly makes all MPC line cards go offline in the MX2000 router while the hardware fault is located on one specific MPC line-card slot. PR1482124
- The vmcore process crashes sometimes along with the mspmand process on MS-MPC and MS-MIC if large-scale traffic flows are processed. PR1482400
- Fragmentation limit and reassembly timeout configuration under services option are missing for SPC3. PR1482968
- On SCBE3, traffic decreases during throughput testing. PR1483100

- The downstream IPv4 packet greater than BR MTU gets dropped in MAP-E. PR1483984
- The traffic rate might not be as expected on the aggregated Ethernet interface after applying a shared-bandwidth policer. PR1484193
- SNMP index in Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in sFlow record data at collector. PR1484322
- On MPC10 line cards, the logical tunnel interface might not work. PR1484751
- In a scaled environment, bbe-smgd might crash when executing the **show system resource-monitor summary** CLI command. PR1484444
- tcpdump core file is generated after initiating **monitor traffic interface** command from CLI. PR1485465
- The **krt-nexthop-ack-timeout** command might not automatically be picked up on restarting the rpd process. PR1485800
- MPC10E line card installed in the FPC slot 4 might drop host outbound traffic. PR1485942
- Kernel core files might be seen if deleting an ifstate. PR1486161
- Kernel crash (vmcore) occurs upon receipt of a malformed IPv6 packet. PR1486948
- On MX Series Virtual Chassis, error logs **ac200_dcfp2_pm_get_info: ifd 0xb3d090a8, info 0x196cc518, ac200_dcfp2 0x0** are generated periodically. PR1487070
- The aftd process might crash. PR1487416
- Incorrect frame length of 132 bytes might be captured in the packet header. PR1487876
- XML is not correctly formatted. PR1488036
- Add support for PSM firmware upgrade on MX2000. PR1488575
- The chassisd process might crash if you execute an SNMP request for a MIC that is a part of an offline FPC. PR1488946
- Daemon might restart due to mishandling of data. PR1489512
- Previous configuration might still take effect after rollback rescue is performed. PR1489575
- With the MX-SPC3 service card, NAT might not be processed on an order as setup. PR1489581
- With a 4-member AMS used in the service set, commit check should fail when a /30 subnet address is used as NAT pool IP. PR1489885
- Support for PSM firmware upgrade on the MX2000 platforms. PR1489939
- Prolonged flow control might occur with MS-MPC or MS-MIC. PR1489942
- Add support for PSM firmware upgrade in utility on the MX2000 platforms. PR1489967
- Syslog error message **Failed to connect to the agentx master agent (/var/agentx/master): Unknown host (/var/agentx/master) (No such file or directory)** is continuously being generated with DNS-sinkhole. PR1490487

- The MPC might crash due to the PHY interface driver issue of MIC in MX2000 and MX10003 platforms. PR1490531
- When a NAT/SFW rule is configured with application-set with multiple applications having different TCP inactivity-timeout values, sessions are not getting TCP inactivity-timeout according to the configured application order. PR1491036
- In an event where BCM SerDes firmware has stopped and not completed, a corresponding alarm is generated. PR1491142
- The ISSU is not supported on the NG-MPC line cards from Junos OS Release 19.4R1. PR1491337
- Multiple deactivation or activation of the security traceoptions along with a single NAPT44 session might crash the flowd process. PR1491540
- MS-MIC goes down after loading some Junos OS releases in an MX Virtual Chassis scenario. PR1491628
- On the MX240, MX480, or MX960 router with SCB3E, swapping the MPC10E line card with the MPC7E line card in the same FPC slot results in fabric errors, which causes system-wide traffic impact. PR1491968
- User-configured MTU might be ignored after the unified ISSU using **request vmhost software in-service-upgrade**. PR1491970
- MX10003 RCB always detect fire temp and shutdown in short time after downgrade. PR1492121
- There is a delay in the LT interfaces on the MPC11E line card coming up after configuring the scaled PS interfaces anchoring to RLT. PR1492330
- A PIC number greater than four will be missing from the SNMP table entPhysicalTable. PR1492996
- The delta PSM firmware upgrade status is incorrectly displayed. PR1493045
- MPC10 and MPC11 line cards might crash if the interface is configured with the firewall filter referencing shared-bandwidth policer. PR1493084
- DHCP subscribers do not come up as expected after deactivating Virtual Chassis port. PR1493699
- The **ptp-clock-global-freq-tracable** leaf value becomes false and does not change to true when the internal lock is in the Acquiring state. PR1493743
- The LSP might not come up in an LSP externally provisioned scenario. PR1494210
- The error message **PFE_ERROR_FAIL_OPERATION: Unable to unbind cos scheduler from physical interface** is seen for the AF interfaces on an FPC when the peer FPC is restarted. PR1494452
- In a node slicing setup, after GRES, the RADIUS interim updates might not carry actual statistics. PR1494637
- B4 devices cannot establish the softwire with AFTR. PR1496211
- The following error messages are generated by Packet Forwarding Engine when the subscribers come up over a pseudowire interface: **PFEIFD: Could not decode media address with length 0**. PR1496265
- Outbound SSH connection flap or memory leak issue might be observed during a push configuration to the ephemeral database with a high rate. PR1497575

- Port numbers logged in the ALG syslog are incorrect. PR1497713
- Subscribers might be disconnected after one of the aggregated Ethernet participating FPCs comes online in a Junos OS node slicing scenario. PR1498024
- SNMP polling does not show correct PSM jnxOperatingState when one of the PSM inputs failed. PR1498538
- The rpd process might crash when multiple VRFs with IFLs link-protection are deleted at a single time. PR1498992
- The commit check might fail when adding a logical interface into a routing instance with the **no-normalization** statement enabled under the routing instances hierarchy. PR1499265
- Heap memory leak might be seen on the MPC10 and MPC11 line cards. PR1499631
- Some of the virtual services might not come up after GRES or rpd restart. PR1499655
- After disabling and enabling the ams0 interfaces, the NAT sessions do not get synchronized back to the current standby SDG. PR1500147
- Inline Junos telemetry interface might report a wrong value for some fields in flow records after enabling nexthop-learning and route churn happens. PR1500179
- The SPC3 card might crash if the SIP ALG is enabled. PR1500355
- The **show services alg conversations** and **show services alg sip-globals** commands are not supported in the USF mode. PR1501051
- On MX2010 and MX2020 routers, the **pem_tiny_power_remaining** message will be continuously logged in chassisd log. PR1501108
- Application ID is not displayed under the NAT/SFW rule configured with application any rule. PR1501109
- VPN traffic gets silently discarded in a cornered Layer 3 VPN scenario. PR1501935
- The chassisd process might become nonresponsive. PR1502118
- On the MPC11 line card, the **show syslog** command in the Packet Forwarding Engine shell might time out. PR1502877
- MACsec delay protection fails to drop or discard delayed MACsec packets. PR1503010
- The packets from a nonexisting source on the GRE or UDP designated tunnel might be accepted. PR1503421
- Configuring the ranges statement for autosensed VLANs might not work on the vMX platforms. PR1503538
- The **show bridge statistics** command output does not display the statistics information for the pseudowire subscriber interfaces. PR1504409
- The gNMI stream does not follow the frequency on the subscription from the collector. PR1504733
- Fan speed might toggle between full and normal on the MX960 router with an enhanced FRU. PR1504867

- S-BFD session might be unable to get up if multiple IP addresses are configured in lo0 interface. PR1505418
- The rpd process might crash in case of a network churn when the telemetry streaming is in progress. PR1505425
- The l2cpd process might crash if the ERP configuration is added or removed, and the l2cpd process is restarted. PR1505710
- GnmiJuniperTelemetryHeader incompatibility is introduced in Junos OS Release 19.3. PR1507999
- The heap memory utilization might increase after extensive subscriber login or logout. PR1508291
- Outbound SSH connection flap or memory leak issues might be observed during a push configuration to the ephemeral database with a high rate. PR1508324
- The ERO update by the controller for branch LSP might cause issues. PR1508412
- False positive TSensor errors are reported on vjunos0. PR1508580
- The host-generated packets might be dropped if the **force-control-packets-on-transit-path** statement is configured. PR1509790
- The disabled QSFP transceiver might fail to be switched on. PR1510994
- Static subscribers are logged out after creating a unit under the demux0 interface. PR1511745
- Memory leak on l2ald might be seen when adding or deleting the routing-instances or bridge-domains configuration. PR1512802
- The wavelength configured through the CLI might not be set on the SFP+-10G-T-DWDM-ZR optics when the optics is used on the MPC7E line card. PR1513321
- Modifying the segment list of the segment routing LSP might not work. PR1513583
- Subscribers might not be able to bind again after performing back-to-back GRES followed by an FPC restart. PR1514154
- The MACsec session might fail to establish if 256-bit cipher suite is configured for MACsec connectivity association assigned to a logical interface. PR1514680
- On the MX2020 and MX2010 routers, the SPMB CPU is elevated when an SFB3 is installed. PR1516287
- Used-Service-Unit of the CCR-U has Output-Bytes counter zero. PR1516728
- The l2ald process crashes during stability test with traffic on a scaled setup. PR1517074
- The MPC7E line card with QSFP installed might get rebooted when the **show mtip-chmac <1|2> registers vty** command is executed. PR1517202
- There might be memory leak in cfmd if both the CFM and inet/IPv4 interfaces are configured. PR1518744
- The vgd process might generate a core file when the OVSDB server restarts. PR1518807
- Traffic loss might happen when an uncorrected (fatal) AER error is detected. PR1519530

- The PADI packets might be dropped when the interface encapsulation VPLS is set along with the accepted protocol configured as PPPoE. PR1523902
- The PSM firmware upgrade must not allow multiple PSM upgrades in parallel to avoid the firmware corruption and support mutliple firmwares for different hardware. PR1524338
- Commit is successful while deactivating CB0 and CB1 interfaces with a running GNF. PR1524766
- According to the OC data model, the openconfig-alarms.yang subscription path must be used as system/alarms/alarm. PR1525180
- Addition and removal of an aggregated Ethernet interface member link might cause the PPPoE subscriber session and traffic to drop. PR1525585
- Error message **Erroneous RPD_DYN_CFG_GET_PROF_NAME_FAILED** is seen during GRES if a RIB interface is configured without a profile. PR1526481
- WAG control route prefix length is observed. PR1526666
- On MX150 routers, physical interface stay up during vmhost halt or power-off. PR1526855
- Family IPv6 does not come up for the L2TP subscriber when additional attributes are not passed in the Framed-IPv6-Route VSA. PR1526934
- Commit error messages come twice while validating the **physical-cores** statement. PR1527322
- The cpcdd process might generate core files after upgrading to Junos OS Release 19.4 and later. PR1527602
- The transit PTP packet might be modified unexpectedly when the packet passes through MPC2E-NG, MPC3E-NG, and MPC5E line cards. PR1527612
- The **commit confirm** command might not roll back the previous configuration when the commit operation fails. PR1527848
- The speed command cannot be configured under the interface hierarchy on an extended port when MX204 or MX10003 router works as an aggregation device. PR1529028
- Non-impacting error message is seen in the message logs: **IFP error> ../../../../../../../../src/pfe/usp/control/applications/interface/ifp.c@3270:(errno=1000) tunnel session add failed**. PR1529224
- The multicast traffic might be dropped due to hash mismatch when there are aggregated Ethernet and ECMP links involved in the multicast tree. PR1529475
- In the subscriber management environment, the RADIUS interim accounting record does not get populated with the subscriber statistics. PR1529602
- SFP-LX10 shows unsupported after unified ISSU upgrade on 3D 20x 1GE(LAN)-E,SFP and 3D 20x 1GE(LAN)-EH,SFP. PR1529844
- PEM 0 always shows as absent or empty even if PEM 0 is present on MX10003 router. PR1531190
- New subscribers might fail to connect due to filter index space exhausted error. PR1531580

- Deletion of the address of the jmgmt0 interface might fail if the shortened version of the CLI command is used. PR1532642
- Some routes might get incorrectly programmed in the forwarding table in the kernel which is no longer present in rpd. PR1534455
- The clear ike statistics with remote gateway does not work. PR1535321
- SNMP MIB walk for jnxSubscriber OIDs returns general error. PR1535754
- Multicast traffic might be sent out through unexpected interfaces with distributed IGMP enabled. PR1536149
- Error message **JAM: Plugin installed for summit_xxx PIC** might be seen when the JAM packages are installed for MX10003 platforms. PR1537389
- The accounting interim-updates for subscriber does not work after GRES and subsequent reboot of the FPCs in a node slicing setup. PR1539474
- Services process mspmand leaks memory in relation to MX telemetry, reporting RLIMIT_DATA exceed. PR1540538
- With hold time configuration, Gigabit Ethernet interfaces remain down on reboot. PR1541382
- Subscriber might not come up on some dynamic VLAN ranges in a subscriber management environment. PR1541796
- Port mirroring with **maximum-packet-length** configuration does not work over a GRE interface. PR1542500
- The PPPoE subscribers might fail to login. PR1551207
- The **show dynamic-profile session client-id** command displays only one IPv6 framed-route information. PR1555476

High Availability (HA) and Resiliency

- Unified ISSU might fail on MX204 and MX10003 Virtual Chassis with an error message. PR1480561

Infrastructure

- If the serial number of the PEM starts with 1F1, the following alarm might be generated: **Minor FPC PEM Temp Sensor Failed**. PR1398128
- Packet counter does not work as expected when using SNMP get. PR1422929
- Unknown MIB OID 1.3.6.1.2.1.47.2.0.30 are referenced in the SNMP trap after upgrading to Junos OS Release 18.4R3.3. PR1508281
- SNMP polling might return an unexpectedly high value for the ifHCOutOctets counter for a physical interface when any jnxDom OID is processed at the same time. PR1508442

Interfaces and Chassis

- Syslog error **scchassisd[ ]: CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments fru_send_msg Global FPC x** is observed after MX Series Virtual Chassis local or global switchover. PR1428254
- Benign **registration is being denied** message might be seen when committing configuration on MX Virtual Chassis. PR1431377
- The MC-LAG configuration-consistency ICL configuration might fail after committing some changes. PR1459201
- The **sonet-options configuration** statement is disabled for the xe interface that works in wan-phy mode. PR1472439
- The interface on MIC3-100G-DWDM might go down after performing an interface flap. PR1475777
- Fail to configure proactive ARP detection. PR1476199
- A stale IP address might be seen after a specific order of configuration changes in a logical systems scenario. PR1477084
- Control logical interface 32767 is not created on the VLAN-tagged physical interface even after removing the VLAN 0 configuration. PR1483395
- Traffic might get dropped as the next hop points to the ICL even though the local MC-LAG is up. PR1486919
- On the MPC6 line cards, the CFM DM two-way verification fails with an invalid timestamp. PR1489196
- The **vrrp-inherit-from** change operation leads to packet loss when traffic is being forwarded to the VIP gateway. PR1489425
- The mgd process might hang up on a crashed dcd commit check process and the dcd might also crash. PR1491363
- The FPC crash might be observed with an inline mode with CFM configured. PR1500048
- Unexpected dual VRRP backup state might occur after performing two subsequent Routing Engine switchovers with track priority-hold-time configured. PR1506747
- Commit failure is observed while deleting all the units under the ps0 interface. PR1514319
- The following error message is observed: **Request failed: OID not increasing: ieee8021CfmStackServiceSelectorType**. PR1517046
- Buffer overflow vulnerability in a device control daemon is observed. PR1519334
- Syslog error **should have at least one member link on a different fpc** might be observed after committing a configuration under interface hierarchy. PR1539719
- The following the commit error is observed while trying to delete unit 1 logical systems interfaces: **ae2.1: Only unit 0 is valid for this encapsulation**. PR1547853

Intrusion Detection and Prevention (IDP)

- When creating the custom IDP signatures that match the raw bytes (hexadecimal), the commit check fails if the administrator configures the depth parameter. PR1506706

J-Web

- Session fixation vulnerability in J-Web. PR1410401
- The httpd process might run with high CPU utilization when J-Web is enabled. PR1483607
- Security vulnerability in J-Web and Web-based (HTTP or HTTPS) services is observed. PR1499280

Juniper Extension Toolkit (JET)

- Behavior change in clients with multiple gRPC channels to same target. PR1492088

Junos Fusion Satellite Software

- Temperature sensor alarm is seen in Junos fusion scenarios. PR1466324

Layer 2 Features

- Connectivity is broken through LAG because of the members configured with hold-time and force-up. PR1481031

Layer 2 Ethernet Services

- For the MX204 router, the vendor ID is set as MX10001 in the factory-default configuration and in the DHCP client messages. PR1488771
- The DHCP subscribers might not come up when DHCP ALQ and VRRP are configured. PR1490907
- JDHCPD memory leak is observed during login or logout test over five days. Memory is more than tripled in this time. PR1491349
- Issues with the DHCPv6 relay processing confirm and reply packets are observed. PR1496220
- The MC-LAG might be down after disabling and then enabling the force-up configuration. PR1500758
- The default-route might not be added to the Juniper device configured as the DHCPv4 client device. PR1504931
- The aggregated Ethernet interface sometimes might not come up after the router is rebooted. PR1505523
- The DHCPv6 lease query is not as expected while verifying the DHCPv6 server statistics. PR1506418
- Receipt of the malformed DHCPv6 packets causes jdhcpd to crash. PR1511782
- The show dhcp relay statistics display DHCPLEASEUNASSIGNED instead of DHCPLEASEUNASSINGED. PR1512239
- The **show dhcpv6 relay statistics** command must display DHCPV6_LEASEQUERY_REPLY instead of DHCPV6_LEASEQUERY_REPL for the messages sent. PR1512246

- The jdhcpd process crashes when a specific DHCPDv6 packet is processed in the DHCPv6 relay configuration. PR1512765
- The DHCP6 lease query is not as expected while verifying the DHCPv6relay statistics. PR1521227
- Memory leak in jdhcpd might be seen if **access-profile** is configured under the **dhcp-relay** or **dhcp-local-server** statement. PR1525052
- Memory leak in the jdhcpd process might be seen if the **access-profile** is configured under the **dhcp-relay** or **dhcp-local-server** statement. PR1525052

MPLS

- The following error is observed on switchover when a vt- interface is in use: **show routing-instances MVPN-1 instance-type vrf; interface vt-0/0/0.1202 { <<<<< multicast; } //snip.**PR1434522
- The RSVP interface bandwidth calculation rounds up. PR1458527
- The rpd process might crash in PCEP for the RSVP-TE scenario. PR1467278
- On the MPC10E and MPC11E line cards, the LDP and BFD sessions are dropped when the fast-lookup-filter has a default term with only accept as action and it is attached to the lo0 interface. PR1474204
- PCC might flood event logs to the controller. PR1476822
- The rpd crash might be seen after back-to-back graceful restart or GRES. PR1485985
- The rpd might crash on restart of the primary Routing Engine or backup Routing Engine when chain-NH has inner and outer labels in the SR-TE scenario. PR1486077
- High CPU utilization for rpd might be seen if RSVP is implemented. PR1490163
- The rpd process might crash when the BGP flaps with FEC 129 VPWS enabled. PR1490952
- BGP session flaps between two directly connected BGP peers because of the incorrect TCP-MSS in use. PR1493431
- The rpd process might crash in a rare condition in the SR-TE scenario. PR1493721
- The rpd core files are generated during unified ISSU. PR1493969
- The rpd process generates core file on the backup Routing Engine. PR1495746
- The rpd process might crash when the SNMP polling is done using the OID jnxMplsTeP2mpTunnelDestTable. PR1497641
- Traffic loss might occur if unified ISSU is performed when P2MP is configured for an LSP. PR1500615
- The CSPF job might get stalled for a new or an existing LSP in a high-scale LSP setup. PR1502993
- The auto-bandwidth feature might not work correctly in an MPLS scenario. PR1504916
- The rpd process might crash with RSVP configured in a rare timing case. PR1505834
- The rpd process might crash when the rpd restarts or GRES switchovers. PR1506062

- Activating or deactivating the LDP-sync under OSPF might cause the LDP neighborship to go down and stay down. PR1509578
- The rpd process might crash after upgrading Junos OS Release 18.1 to a later release. PR1517018
- The SNMP trap is sent with the incorrect OID jnxSpSvcSetZoneEntered. PR1517667
- The LDP session-group might throw a commit error and flap. PR1521698
- The inter-domain LSP with a loose next hops path might get stuck in down state. PR1524736
- LDP routes might be deleted from MPLS routing table after Routing Engine switchover. PR1527197
- The **ping mpls rsvp** command does not take into account the lower MTU in the path. PR1530382
- The rpd process might crash when the LDP route with indirect next hop is deleted on the aggregated Ethernet interface. PR1538124
- Performing commit might trigger externally provisioned LSP MBB mechanism. PR1546824

Network Address Translation (NAT)

- Improve the max ENODE connections for one persistent NAT binding from 8 to 32. PR1532249

Network Management and Monitoring

- The SNMPv3 informs might not work properly after rebooting. PR1497841

Platform and Infrastructure

- The **core.vmxt.mpc0** is seen at **5 0x096327d5 in l2alm_sync_entry_in_pfes (context=0xd92e7b28, sync_info=0xd92e7a78) at ../../../../../src/pfe/common/applications/l2alm/l2alm_common_hw_api.c:1727**. PR1430440
- Traffic loss might be seen in case of Ethernet frame padding with VLAN. PR1452261
- Traffic from an IRB interface toward an LSI interface gets dropped with adaptive or per-packet load balancing. PR1458825
- On the MX204 router, GRE with sampling causes the following Packet Forwarding Engine error: **MQSS(0): MALLOC: Underflow error during reference count read - Overflow 1, Underflow 1, HMCIF 0, Address 0x8d62e0**. PR1463718
- VXLAN packet might be discarded with flow caching enabled on MX150 and vMX. PR1466470
- SSH login might hang and the TACACS+ server closes the connection without sending any authentication failure response. PR1478959
- Traffic disruption for an MoFRR protected multicast flow in an NG-MVPN hot root standby FRR scenario. PR1478981
- Some error logs might be reported every 2 minutes due to SRAM single bit ECC error which is a transient hardware issue on MX Series with MPCs/MICs with queuing chip. PR1479240
- XQCHIP xqchip_drop_get_q_length severity of parity error moved from major to minor. PR1481558

- The **show system buffer** command displays all zeros in the MX104 chassis. PR1484689
- MAC learning under bridge domain stops after the MC-LAG interface flaps. PR1488251
- Normalize PPE thread timeout settings across platforms with different clock speeds. PR1490761
- MAC malformation might happen in a rare scenario under MX Series Virtual Chassis setup. PR1491091
- In a node slicing setup, MPLS TTL might be set to zero when the packet goes through AF interface configured with CCC family. PR1492639
- Traceroute monitor with MTR version v.69 shows a false 10 percent loss. PR1493824
- Packets get dropped when the next hop is an IRB-over-LT interface. PR1494594
- The Routing Engine might crash when a large number of next hops are quickly deleted and added again in a large ARP or ND scaled scenario. PR1496429
- Traffic to VRRP virtual IP or MAC addresses might be dropped when ingress queuing is enabled. PR1501014
- Python or SLAX script might not be executed. PR1501746
- MAC learning request throttling mechanism could not work properly in a scale setup. PR1501758
- Arbitrary code execution vulnerability in Telnet server. PR1502386
- Traffic originated from another subnet is sent out with 0x8100 instead of 0x88a8. PR1502867
- MPCs might crash when there is a change on routes learned on the IRB interface configured in the VPLS or EVPN instances. PR1503947
- Traffic loss might be seen in certain conditions under an MC-LAG setup. PR1505465
- The kernel might crash causing the router or the Routing Engine to reboot when performing virtual IP related change. PR1511833
- During route table object fetch failure, the FPC might crash. PR1513509
- The output of the **show jnh qmon queues-sensor stats 0** command has no content. PR1514881
- VPLS connection might be stuck in primary fail status when a dynamic profile is used on the VPLS pseudowire logical interface. PR1516418
- The configured scheduler map is not applied on the ms- interface if the service PIC is in the Offline state during commit. PR1523881
- TWAMP interoperability issue between Junos OS releases are observed. PR1533025
- NH DWORD memory leak observed in the Packet Forwarding Engine with ARP churn bound to IRB interface, part of EVPN-MPLS routing instance. PR1533857
- Subscribers are not coming up on PS interface. PR1536043

- The rmopd process memory leak might be seen if TWAMP client is configured. PR1541808
- MX Series with MPCs/MICs might crash when the underlying Layer 2 interface for ARP over IRB interface is changed from physical interface to LSI interface. PR1542211

Routing Policy and Firewall Filters

- The router ID from the martian address range cannot be committed even if the range is allowed by the configuration. PR1480393
- The policy configuration might be mismatched between rpd and mgd when the **deactivate policy-options prefix-list** is involved in configuration sequence. PR1523891

Routing Protocols

- The BGP session might become nonresponsive with high BGP OutQ value after GRES on both sides. PR1323306
- When configuring an alternate incoming interface for a PIM RPF check using rpf-selection, the additional groups outside the configured range might switch to the alternate incoming interface. PR1443056
- IS-IS TI-LFA traffic convergence time is more than 50 ms for IPv4 and IPv6 traffic. PR1458791
- Adjacency SID might be missed and not be advertised to peer/controller/BMP monitor in BGP-LS NLRI. PR1473362
- The rpd process crashes due to specific BGP UPDATE packets. PR1481641
- Multicast traffic loss might be seen in certain conditions while enabling the IGMP snooping under an EVPN-VXLAN ERB scenario. PR1481987
- The rpd process might crash when deactivating logical systems. PR1482112
- The BGP multipath traffic might not fully load-balance for a while after adding a new path for load sharing. PR1482209
- The output of the **show isis interface detail** command might be incorrect if **wide-metrics-only** is enabled for IS-IS and the ASCII representation of the metric in decimal is more than 6 characters. PR1482983
- RIPv2 might malfunction when changing the interface type from P2MP to broadcast. PR1483181
- The rpd process crashes if the same neighbor is set in different RIP groups. PR1485009
- There might be rpd process memory leak in a certain looped MSDP scenario. PR1485206
- The BGP-LU routes do not have the label when BGP sharding is used. PR1485422
- Removal of the BGP and rib-sharding configuration might cause the routing protocols to become unresponsive. PR1485720
- Layer 3 VPN RR with the **family route-target** and **no-client-reflect** statements does not work as expected. PR1485977
- Traffic loss might be seen while performing GRES in an MPLS setup. PR1486657

- The rpd crashes if BGP LLGR with RIB sharding and traceoptions for graceful-restart are configured. PR1486703
- The rpd might crash with BGP RPKI enabled in a race condition PR1487486
- The rpd might crash when you perform GRES with MSDP configured. PR1487636
- High CPU utilization might be observed when the outgoing BGP updates are sent slowly. PR1487691
- The rpd process might generate core files after **always-compare-med** is configured for BGP path selection. PR1487893
- The BGP RIB sharding feature cannot be run on a system with a single CPU. PR1488357
- The rpd crashes when OSPF neighbors are reset. PR1489637
- Ppmd core file is generated after MS-MPC restart. PR1490918
- The BGP route target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743
- A core file is generated in **krt_mcnh_update_rpf_info()** when TI-LFA is used with MOFRR. PR1493259
- The rpd process generates core files at **rt_nh_resolve_add_gen in ../../../../../../../../src/junos/usr.sbin/rpd/lib/rt/rt_resolve_ind.c:** with the EVPN DHCP configurations. PR1494005
- The static route in inet6.0 or inet6.3 RIB might fail to be deleted. PR1495477
- Receipt of certain genuine BGP packets from any BGP speaker causes the rpd process to crash. PR1497721
- The route entries might be unstable after being imported into inet6.x RIB through a RIB group. PR1498377
- The rpd process might crash if the import policy is changed to accept more routes that exceed the teardown function threshold. PR1499977
- The rpd process might crash in a multicast scenario with BGP configured. PR1501722
- The rpd process might crash while processing a specific BGP packet. PR1502327
- On MX series Virtual Chassis, when you run the **show bgp neighbors** command, change in the x-path output for the value input-updates is observed. PR1504399
- BGP might not advertise routes to peers after a peer flap. PR1507195
- The rpd crash might be seen on a new primary Routing Engine if switchover happens with massive routing instances deletion. PR1507638
- The rpd process might crash due to RIP updates being sent on an interface in the down state. PR1508814
- The rpd process might crash on the backup Routing Engine if BGP (standby) receives a route from the peer, which is rejected due to an invalid target community. PR1508888
- The rpd might report 100 percent CPU usage with the BGP route damping enabled. PR1514635

- IS-IS segment routing routes might not be updated to reflect the change in the SRMS advertisements. PR1514867
- The rpd process might crash after deleting and then adding a BGP neighbor. PR1517498
- The rpd process might crash if there is a huge number of SA messages in an MSDP scenario. PR1517910
- Tag matching in the VRF policy does not work properly when the **independent-domain** option is configured. PR1518056
- Need BGP-LS NLRI handling improvements for BGP-LS ID TLV. PR1521258
- The VRF label is not assigned at ASBR when the inter AS is implemented. PR1523896
- VRF label is not assigned at ASBR when inter AS is implemented. PR1523896
- The IS-IS LSP database synchronization issue might be seen while using the flood-group feature. PR1526447
- Transit labels for Layer 3 VPN routes pushed momentary to mpls.0 table. PR1532414
- Configuring **then next-hop** and **then reject** on a route policy for the same route might cause rpd crash. PR1538491
- After move peer out of protection group, path protection is not removed from the PE router and multipath route is still present. PR1538956
- The rpd process generates the core file at **gp_rtarget_tsi_update,bgp_rtarget_flash_rt,bgp_rtarget_flash**. PR1541768

Services Applications

- The fpc process might crash with the npc core file if the service interface is configured under a service set in USF mode. PR1502527
- The output of the **show services l2tp tunnel extensive** command does not show the configured session limit. PR1503436
- Destination lockout functionality does not work at the tunnel session level when CDN code is received. PR1532750

Subscriber Access Management

- Subscriber accounting messages retransmissions exist even after configuring accounting retry 0. PR1405855
- The following syslog message is observed: **pfe_tcp_listener_open_timeout: Peer info msg not received from addr: 0x6000080. Socket 0xfffff804ad23c2e0 closed**. PR1474687
- The delete request of a specified service session through CoA could fail. PR1479486
- NAS-Port-ID includes a subinterface in the RADIUS messages for the aggregated Ethernet interface. PR1484351
- The authd log events might not be sent to the syslog host when **destination-override** is used. PR1489339

- The LTS incorrectly sends the access-request with the Tunnel-Assignment-ID, which is not compliant with RFC 2868. PR1502274
- CCR-T does not contain the usage monitoring information. PR1517507
- The **show network-access aaa subscribers statistics username "<>"** command fails to fetch the subscriber-specific AAA statistics information if the username of the subscriber contains space. PR1518016

User Interface and Configuration

- The version information under the configuration changes from Junos OS Release 19.1 onward. PR1457602

VPNs

- Traffic loss is observed while verifying a multicast route with VT interface for VPNA. PR1460480
- In an MVPN scenario, the LSP might stay down on removing all VT interfaces from a single-hop egress. PR1474830
- The Layer 2 circuit neighbor might become nonresponsive in the Ready state at one end of the MG-LAG peer. PR1498040
- The rpd process might crash in certain conditions after deleting the Layer 2 circuit configuration. PR1502003
- The MPLS label manager might allow configuration of a duplicated VPLS static label. PR1503282
- The rpd process might crash after removing the last configured interface under the Layer 2 circuit neighbor. PR1511783
- The rpd process might crash when deleting the Layer 2 circuit configuration in a specific sequence. PR1512834

Resolved Issues: 20.1R1

Application Layer Gateways

- SIP messages that needs to be fragmented might get dropped by the SIP ALG. PR1475031

Authentication and Access Control

- The LLDP packets might get discarded on all Junos OS platforms. PR1464553

Class of Service (CoS)

- The MX Series generated OAM/CFM LTR messages are sent with a different priority than the incoming OAM/CFM LTM messages. PR1466473
- Unexpected traffic loss might be discovered in certain conditions under in a Junos fusion scenario. PR1472083

- The MX10008 and MX100016 routers might generate cosd core files after executing the **commit/commit check** command if the **policy-map** configuration is set. PR1475508

EVPN

- Traffic received from VTEP is dropped if the VNI value used for type-5 routes is greater than 65,535. PR1461860
- Rpd might crash with the EVPN-related configuration changes in a static VXLAN to MPLS stitching scenario. PR1467309

Forwarding and Sampling

- Traffic errors do not get policed as expected after being locally switched for VLAN 100 and 101, while verifying the selective local-switching functionality with 4000 VLANs. PR1436343
- The pfed might crash and not be able to come up on the PTX Series or TVP based platforms. PR1452363
- The following syslog error messages are seen: **pfed: rtslib: ERROR received async message with no handler: 28.** PR1458008
- The following false warning message is seen on commit (commit check) after upgrading to Junos OS Release 19.2R2-S1.4: **warning: vxlan-overlay-load-balance configuration for forwarding options has been changed.** PR1459833
- On an MX Series router, the following logs are seen: **L2ALD_MAC_IP_LIMIT_REACHED_IF: Limit on learned MAC+IP bindings reached for .local.1048605; current count is 1024**. PR1462642
- Type 1 ESI/ or AD routes are not generated locally on EVPN PE devices in all-active mode. PR1464778
- On the MX10008 and MX10016 routers, policer bandwidth-limit cannot be set higher than 100-Gigabit Ethernet. PR1465093
- An output bandwidth-percent policer with logical-bandwidth-policer applied to an aggregated Ethernet bundle along with an output-traffic-control-profile has incorrect effective policing rate. PR1466698
- Traffic might be forwarded into the default queue instead of the right queue when the VPLS traffic has three or more VLAN tags with VLAN priority 5. PR1473093
- The filter might not be installed if the **policy-map xx** is present under the filter. PR1478964

General Routing

- The severity of the following error is reduced from fatal to major: **XR2CHIP_ASIC_JGCI_FATAL_CRC_ERROR.** PR1390333
- On the MX240, MX480, or MX960 router with SCB3E, swapping MPC10E line card with MPC7E line card in the same FPC slot results in fabric errors, which causes system-wide traffic impact. PR1491968
- After restarting the routing or rpd process, sometimes the sensor statistics is not reset. After the rpd process restarts, the sensor do not reset and traffic statistics increases on the existing value. PR1458107

- A newly added LAG member interface might forward traffic even though its micro BFD session is down. PR1474300
- In a configuration mode, when you ask for command completion help for the **co-ordinate configuration** statement at the **[edit protocols lldp-med interface location]** hierarchy level, you see that the word value is misspelled in the help text. PR1486327
- On the MX104 platform with any 2-port license installed on the 10-Gigabits Ethernet interfaces and phy-timestamping enabled in PTP, PTP might not work. PR1421811
- Default configuration does not create any logical interfaces and LLDP cannot discover neighbor for those interfaces which logical interface is not configured explicitly in the Junos OS configuration. PR1436327
- The failover time for the LACP link protection might be more than 2 seconds on the MPC11E line card. PR1464652
- The following constant messages flooding in log is observed: **summit_pic_port_profile_isvalid: VALID Port profile.** PR1464879
- The **high-cos-queue-threshold** range is changed to [uint 0 .. 90;]. PR1390424
- NAPT66 pool split is not supported with AMS; thus commit must fail with IPv6 pool in AMS. PR1396634
- The non existent subscribers might appear in the **show system resource-monitor subscribers-limit chassis extensive** output. PR1409767
- Changing CAK and CKN multiple times within a short interval (around 5 minutes) sometimes show the security MACsec connection's inbound and outbound channel display with more than one active AN. But on the Packet Forwarding Engine hardware side, the correct AN and SAK is programmed and MKA protocol from both ends transmits the correct and latest AN on each hello packet. You should not see any traffic drop due to this display issue. PR1418448
- Certain JNP10008-SF and JNP10016-SF Switch Interface Boards (SIBs) manufactured between July 2018 and March 2019 might have incorrect core voltage setting. PR1420864
- The jnxFruState shows value as 10 for Routing Engine instead of 6 in response to .1.3.6.1.4.1.2636.3.1.15.1.8.9.1.0.0. PR1420906
- Ports might get incorrectly channelized if they are already of 10-Gigabit Ethernet and they are channelized to 10-Gigabit Ethernet again. PR1423496
- Observing NPC core files at **trinity_rtt_hw_bulk_helper**, **trinity_rt_delete**, rt_entry_delete_msg_proc (rt_params=0x48803bd8) at **../../../../../../../../src/pfe/common/applications/route/hal/rt_entry.c:5210**. PR1427825
- The following syslog error message is observed: **Err] dfw_abstract_issu_stats_counters_restore:2222 Failed to find Index = 4613734? during ISSU with 19.3I-20190409_dev_common.0.2212.** PR1429879
- The routers that are configured with the protect core file might send IPfix sampling packets with the incorrect next-hop information. PR1430244

- The l2cpd process might crash and generate a core file when the interfaces flap. PR1431355
- MicroBFD 3x100ms flap is observed upon inserting a QSFP in another port. PR1435221
- ZF interrupts for out-of-range destination Packet Forwarding Engine INTR for Gnt is observed when the MPC6 or MPC9 line card is brought up. PR1436148
- ISSU fails from the legacy Junos OS Release 19.1R1 images. PR1438144
- Incorrect values are observed in the **JUNIPER-TIMING-NOTFNS-MIB** table. PR1439025
- The ports of the EX devices might stay in the **Up** state even if the EX4600 or QFX5100 lines of switches is rebooted. PR1441035
- The interface might go into the **Down** state after the FPC restarts with the PTP configuration enabled. PR1442665
- The BGP session fails to establish when you use the firewall filter to de-encapsulate BGP packets from the GRE tunnel. PR1443238
- System reboot is required when GRES is enabled or disabled with the **mobile-edge** configuration. PR1444406
- Irregular traffic drop might be seen when the traffic is ingress from MPC3E and egress to MPC10E. PR1445649
- When you use a converged CPCD, an MX Series router rewrites the HTTPS request with the destination-port 80. PR1446085
- When switchover happens with an MX Series router with service interface that has NAT and GR configuration, the static route for NAT never comes up. PR1446267
- DT_BNG: bbe-smgd generates core file on the backup Routing Engine in **bbe_ifd_add_vlan (ifd=0x8c3e835, ifl=0xcaf59f18)** at **../../../../../../src/junos/usr.sbin/bbe-svcs/smd/infra/bbe_ifd.c:6374**. PR1447493
- IPv6 throughput numbers for NAT with HTTP traffic are not at par with IPv4. PR1449435
- Changing the hostname triggers the LSP on-change notification and not the adjacency on-change notification. PR1449837
- On the MPC10E line card, dcd is unable to clean stale the mt- logical interfaces while reloading rosen configuration on the DUT. PR1450953
- When you use the Standard_D5_v2, which has 16 vCPUs and 56 GB of memory, the deployment fails. PR1450975
- JNP10000-LC2101 FPC generates **Voltage Tolerance Exceeded** major alarm for each IP 2V5 sensor. PR1451011
- Main chassisd thread at the JNS GNF might stall upon the GNF SNMP polling for hardware-related OIDs. PR1451215
- Need to add support for drop flows when the packet drops. PR1451921

- On the MX10000 and PTX10000 lines of routers with Routing Engine redundancy configuration enabled, the firmware upgrade for PSU (JNP10000-AC2) and JNP10000-DC2) might fail due to lcmd being disabled by the firmware upgrade command. PR1452324
- Sensord core file might be seen when the script runs on MPC10E line card. PR1452976
- On an MPC10E line card, inconsistency between AFT and non-AFT line cards occurs while displaying ldp p2mp traffic-statistics on the bud node. PR1453130
- Add the **syslog** configuration command to the stateful firewall rule then condition. PR1453502
- On an MX10003 device, alarms are not sent to syslog. PR1453533
- The VMX might work abnormally in a large topology. PR1453967
- The 100-Gigabit Ethernet interfaces might not come up again after going down on MPC3E-NG. PR1454595
- When the scale configurations are applied, chassisd CLI command might delay response or might time out for 10 minutes. PR1454638
- On the line card, interface damping is not supported. PR1455152
- The smihelperd process is not initialized when Junos OS is upgraded on PPC-based platforms. PR1455667
- Multiple daemons might crash on committing configuration changes related to groups. PR1455960
- Along with the 4x1GE feature using the QSFP28 optics, continuous logging in the chassisd file is observed when speed 1-Gigabit Ethernet is configured with **pic_get_nports_inst** and **ch_fru_db_key**. PR1456253
- On the line card, need to add the support of optics-options low light. PR1456894
- The bbe-statsd process might continuously crash if any parameter is set to 0 in the **mx_large.xml** file. PR1457257
- On the JSU package installed for lcmd, the daemon might not restart the daemon with the new daemon package. PR1457304
- The chassisd process and all FPCs might restart after Routing Engine switchover. PR1457657
- After more than 2 million multicast subscribers are activated without performing GRES or bbe-smgd restart, further multicast subscribers might be unable to log in. PR1458419
- Traffic silently discards or MPC crashes on the MPC10E line card during the change of the firewall filter terms. PR1458499
- If you use the dynamic VoIP VLAN assignment, the correct VoIP VLAN information in LLDP-MED packets might not be sent after you commit. PR1458559
- The FPC X major errors alarm might be raised after committing the PTP configuration change. PR1458581
- The rpd crash might be seen if the BGP route is resolved over the same prefix protocol next hop in the inet.3 table that has both the RSVP and LDP routes. PR1458595

- The traffic might be stuck on MS-MPC or MS-MIC with sessions receiving a huge number of affinity packets. PR1459306
- The following error message might be seen after the chassisd restarts: **create_pseudos: unable to create interface device for pip0 (File exists).** PR1459373
- The **show ancp subscriber access-aggregation-circuit-id < access aggregation circuit ID>** command displays incomplete output. PR1459386
- Telemetry streaming of mandatory TLV ttl learned from LLDP neighbor is missing. PR1459441
- The traffic might be silently dropped or discarded during the link recovery in an open Ethernet access ring with ERPS configured. PR1459446
- Inline S-BFD packets are dropped on MPC6E MIC1/PIC1 ports: 0-11. PR1459529
- In an MC-LAG scenario, the traffic destined to VRRP-virtual MAC gets dropped. PR1459692
- After the DRD auto-recovery, the traffic is silently dropped or discarded during interface flaps. PR1459698
- Configuration change might not be applied if the Ephemeral database is used. PR1459839
- Initial synchronization for the **OpenConfig** event sensors are streamed only from producers supporting event paths. PR1459927
- On the line card, interface flaps multiple times after an administrator disables or enables at the side or when an optical module is plugged into. PR1459942
- In a subscriber management environment, subscriber statistics reported by CLI commands and RADIUS can be broken if ISSU is performed from any Junos OS Release earlier than 18.4 to 18.4 or later. PR1459961
- The PPTP does not work with destination NAT. PR1460027
- If **vlan-offload** is configured on the VMX platform, **input-vlan-map** might not work. PR1460544
- Support of **del_path** for the LLDP neighbor changes at various levels. PR1460621
- When you receive IPv6 over IPv4 IBGP session, the IPv6 prefix is hidden. PR1460786
- The PTP function might consume the kernel CPU for a long time. PR1461031
- Explicit Deletion Notification (del_path) are not received when the LLDP neighbor is lost as result of disabling the local interface on the DuT through CLI (gNMI). PR1461236
- The bbe-smgd generates a core file when all RADIUS servers are unreachable. PR1461340
- Traffic might be impacted due to fabric hardening being stuck. PR1461356
- The traffic might not be forwarded when it is received from the circuit cross-connect interface. PR1461532
- On the MPC10E line card, more output packets are seen than expected when the ping function is performed. PR1461593

- In an EVPN scenario, memory leak might be observed when **proxy-macip-advertisement** is configured. PR1461677
- The repd generates a core file during system startup. PR1461796
- During the BBE statistics collection and management process, issues with the bbe-statsd memory on the backup Routing Engine occur. PR1461821
- JET RIB API RouteRemove and RouteRemoveMatching RPCs do not work as the first RIB API call. PR1461974
- The rpd might crash after committing the **dynamic-tunnel-anchor-pfe** command. PR1461980
- The rpd process might crash if the **show v4ov6-tunnels information anti-spoof-ip** command is executed. PR1462047
- The following error message appears when both the DIP switches and power switch are turned off: **CHASSISD_SNMP_TRAP6: SNMP trap generated: Power Supply failed.** PR1462065
- The flow stuck and flowd watchdog generate core files while trying to ping the DNS server 8.8.8.8 on the internet through DUT configured with NAPT44. PR1462277
- Traffic drops over the aggregated Ethernet interfaces configured with Virtual Router Redundancy Protocol (VRRP). PR1462310
- On an MX204 router, the RADIUS interim accounting statistics are not populated. PR1462325
- The EA WAN SerDes gets into the **Stuck** state that leads to continuous **DFE tuning timeout** errors and causes the link to stay down. PR1463015
- The vty remote MAC addresses are not learned with correct age if vty is from a line card without Juniper Trio 5 silicon. PR1463040
- MAC-learning is broken for vlan-id all scenario. PR1463078
- The Routing Engine switchover might not be triggered when the master CB clock fails. PR1463169
- MVPN traffic might be dropped after performing switchover. PR1463302
- The subscribers might not pass traffic after making some changes to the dynamic-profiles filter. PR1463420
- RPC ALG causes MSPMAND to generate core files when an MX Series router is used as a stateful firewall with the MS-MIC or MS-MPC service cards. PR1464020
- The IPoE subscriber route installation might fail. PR1464344
- Observing **bbe-smgd-core (0x000000000088488c** in **bbe_autoconf_delete_vlan_session_only (session_id=918)** at **../../../../../../src/junos/usr.sbin/bbe-svcs/smd/plugins/autoconf/bbe_autoconf_plugin.c:3115)**. PR1464371
- The PPP IPv6CP might fail if the **routing-services** command is enabled. PR1464415

- The CPU utilization on mgd daemon might get stuck at 100 percent after the netconf session is interrupted by flapping interface. PR1464439
- The MS-MIC might not work when it is used on a specific MPC. PR1464477
- The **show task memory detail** command shows incorrect cookie information. PR1464659
- The PPPoE session goes in to the **Terminated** state and the accounting stops for the session that is delayed. PR1464804
- MPC5E or MPC6E might crash due to internal thread hogging of the CPU. PR1464820
- The end in front of NAT also sends NATT keep alive packets. PR1464864
- Commit script does not apply changes in the private mode unless a commit full is performed. PR1465171
- The jdhcpd might consume high CPU and no further subscribers can be brought up if more than 4000 dhcp-relay clients are present in the MAC-MOVE scenario. PR1465277
- The physical interface of aggregated Ethernet might take time to come up after disabling or enabling the interface. PR1465302
- Bandwidth percent with shaping rate does not work on an aggregated Ethernet interface after deactivating and activating the class of service. PR1465766
- ICMP error messages does not appear even enabling the **enable-asymmetric-traffic-processing** statement. PR1466135
- The PPPoE subscribers get stuck due to the PPPoE inline keepalives that do not work properly. PR1467125
- Layer 2 wholesale does not forward all the client requests with stacked VLAN. PR1467468
- Hot-swapping between MPC11E and legacy MPC9, MPC8, or MPC6 is not supported. PR1467725
- The process rpd might crash after making several changes to the flow-spec routes. PR1467838
- Crypto code might cause high CPU utilization. PR1467874
- You might observe the following error message: **the user-ad-authentication subsystem is not responding to management requests**. PR1467991
- The **satellite-management** commands are not available. PR1467997
- Benign logs might show in Junos OS Release 19.3R2 when switching between configurations using **load-override** with GRES and **commit-synchronize**. PR1468234
- Optics measurements might not be streamed for the interfaces of a PIC over JTI. PR1468435
- The process rpd crash might be seen if the BGP sharing is enabled. PR1468676
- The **Inner-list** functionality with dual tag does not work. Traffic gets dropped at the ingress port. PR1469396
- The tcp-log connections fail to reconnect and get stuck in the **Reconnect-In-Progress** state. PR1469575
- Memory leak on Layer 2 cpd process causes Layer 2 cpd to crash. PR1469635

- A hierarchical-scheduler should not be configured on a ps- interface. PR1470049
- On the MPC11E line card, some of the 10-Gigabit Ethernet interface states might not get cleaned up correctly when performing GRES with invalid profile configuration. PR1470153
- On MPC-11E interfaces, certain configuration steps might cause traffic to not get policed properly. PR1470629
- The SNMP interface-mib stops working for the PPPoE clients. PR1470664
- On MPC11E, PIC online event does not generate SNMP trap when PIC goes through offline to online transition. PR1470796
- Unable to setup 26M sessions (NAPT44) at 900,000pps per second. PR1470833
- On rare occasions, the router might send out one extra URR quota value for a bearer. PR1470890
- Sudden FPC shutdown due to hardware failure or ungraceful removal of line card might cause major alarms on other FPCs in the system. PR1471372
- In the cRPD platform, license violations are captured as nagging log messages and no alarm is raised. PR1471455
- The clksyncd crash might be seen when PTP over an aggregated Ethernet interface is configured on the MX104 platform. PR1471466
- Phase or frequency synchronization might not work correctly when PTP is configured in the hybrid mode. PR1471502
- MTU errors count captured in the **show pfe statistics traffic** does not match exactly to the actual count of the frames dropped. PR1471554
- On the MX10008 and MX10016 line cards, the ARP suppression (default enabled) in EVPN does not work. PR1471679
- PCC tries to send a report to PCE but the connection between PCC and PCE is not in the **Up** state especially in the case of MBB in PCE provisioned or controlled LSP. PR1472051
- On multicore next-generation Routing Engines on the MX960, MX240, and MX480 routers with USF mode enabled and USF-based services configuration, the subsequent Junos vmhost upgrade fails with an error message. PR1472287
- Chassis alarm on BSYS might be observed : **RE0 to one or many FPCs is via em1: Backup RE.** PR1472313
- Service accounting statistics do not get updated after changes are made to the firewall filters. PR1472334
- The kernel might crash and vmcore might be observed after the configuration change is committed. PR1472519
- Performing back-to-back rpd restarts might cause rpd to crash. PR1472643
- Active error counts do not increase for I2C in the synchronization cards. PR1472660

- On the MX Series devices, if the **reauthenticate lease-renewal** statement is enabled for DHCP, when the DHCP authentication and re-authenticate lease-renewal occurs, the SDB might go down very frequently. PR1473063
- Drops counter does not increment for the aggregated Ethernet even after the member link shows the drops. PR1473665
- Ingress multicast replication does not work with the GRES configuration. PR1474094
- An MPC11 crash might occur on the MX2000 platform using multi dimensional advanced scale configuration that has inline keep alive sessions. PR1474160
- MX10000 QSA adapter lane 0 port goes in the **Down** state when you disable one of the other lanes. PR1474231
- With URR enabled, the URR reports cause memory leak. Eventually, the heap memory gets exhausted. PR1474306
- The **show services sessions** and **show services sessions extensive output** commands do not display the member interface of the AMS where the session got landed. They display only the AMS interface name. PR1474313
- When traffic loss is observed on a 100-Gigabit Ethernet logical interface, the MACsec sessions are up and live. PR1474714
- The **request system power-off** and **request system halt** commands might not work correctly. PR1474985
- The clksyncd generates core files after GRES. PR1474987
- SFW rule configuration deletion might lead to memory leakage. PR1475220
- The Radius accounting updates of the service session have incorrect statistic data . PR1475729
- Dark window size is more than expected and 31.0872721524375 seconds of traffic loss is observed. PR1476505
- The bbe-mibd might crash on the MX Series platform in a subscriber environment. PR1476596
- The MX Series router acting as LNS does not get to program the Packet Forwarding Engine with l2tp services, which causes forwarding issues for the l2tp subscribers. PR1476786
- Traffic loss might be seen in the SAEGW scenario after the daemon restarts or after the GRES operation. PR1477461
- IKE version 2 tunnel flaps with DPD occur if initiator is not behind NAT. PR1477483
- The Packet Forwarding Engine might be disabled due to major errors on MPC2E-NG, MPC3E-NG, MPC5, MPC6, MPC7, MPC8, and MPC9 line cards. PR1478028
- The **show evpn statistics instance** command gets stuck on the multihomed scenario. PR1478157
- At scale log ins of both the default and dedicated bearers might require retries from the control plane. PR1478191
- FPC memory leak might happen after executing the **show pfe route** command. PR1478279

- [firewall] [filter_installation] Output chain filter counters are not correct. PR1478358
- The core files are generated at **cassis_alloc_list_timed_free in cassis_free_thread_entry**. PR1478392
- The protocol MTU might not be changed on the lt- interface from the default value. PR1478822
- The TCP-log sessions might be in the **Established** state but no logs get sent out to the syslog server. PR1478972
- The rpd process might crash when executing the **show route protocol l2-learned-host-routing** or **show route protocol rift** command on a router. PR1481953
- The MX204 router reboots when the PPPoE client starts to log in and no core files are generated. PR1482431
- Packet loss might be observed after the device reboots or l2ald restarts in an EVPN-MPLS scenario. PR1484468
- UID might not be released properly in some scenarios after the service session deactivation. PR1188434
- The **show subscriber extensive** command incorrectly displays DNS address provided to the DHCP clients. PR1457949
- PPP IPv6 NCP fails to negotiate during the PPP login. PR1468414
- DHCP relay with forward-only fails to send OFFER when the client is terminated on the lt-0/0/0.2 logical tunnel interface. PR1471161
- Dynamic-profile for VPLS-PW pseudowire incorrectly reports the Dynamic Static Subscriber Base Feature license alarm. PR1473412
- DHCP-server RADIUS given mask is being reversed. PR1474097

Infrastructure

- The kernel crashes during the removal of the mounted USB when a file is being copied to it. PR1425608
- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1. PR1462986
- The scheduled tasks might not be executed if the cron daemon goes down without restarting automatically. PR1463802

Interfaces and Chassis

- Restarting chassisd with GRES disabled might cause FPC to restart and some demux interfaces to be deleted. PR1337069
- When the logical interface is associated to a routing-instance inside a LR, the logical interface is removed from the routing-instance and the logical interface is not added to the default routing instance. PR1444131
- Continuous VRRP state transition (VRRP master or backup flaps) is observed when one device drops the VRRP packets. PR1446390
- Interface descriptions might be missing under the logical systems CLI. PR1449673

- Mismatched MTU value causes the RLT interface to flap. PR1457460
- The EOAM CFM primary-vid functionality does not work if the **enhanced-cfm-mode** is enabled. PR1465608
- vrrpv3mibs does not work on the QFX platform to poll the VRRPv6 related objects. PR1467649
- The voltage high alarm might not be cleared when voltage level comes back to normal for MIC on MPC5. PR1467712
- When you configure ESI on a physical interface, the traffic drops when you disable the logical interface under the physical interface. PR1467855
- When dynamic DHCP sessions exist in the device and if multiple commits in parallel are performed, the commit might become nonresponsive. PR1470622
- Commit error was not thrown when the member link was added to multiple aggregation groups with different interface specific options. PR1475634
- When the addition and the deletion of an logical interface (both logical interfaces with the same VLAN ID) is performed in a single commit configuration, the check fails with the following error message: **duplicate VLAN-ID.** PR1477060
- MC-AE interface might be shown as an unknown status when you add the sub interface as part of the VLAN on the peer MC-AE node. PR1479012
- For ATM interfaces configuration, if any logical interface has the **allow-any-vci** configuration, then the commit operation might fail. PR1479153

Junos Fusion Enterprise

- Loop detection might not work on the extended ports in the Junos fusion scenarios. PR1460209

Layer 2 Ethernet Services

- The jdhcpd process might go into infinite loop and cause CPU full utilization. PR1442222
- DHCP subscriber might not come online after the router reboots. PR1458150
- On the MX2010 and MX2020 lines of routers, no alarm is generated when FPC is connected to the master Routing Engine through the backup Routing Engine. PR1461387
- The metric does not change when configured under DHCP. PR1461571
- Member links state might be asynchronized on a connection between PE and CE devices in the EVPN A/A scenario. PR1463791
- The ISSU might fail during the subscriber in-flight login. PR1465964
- Telemetry data for **relay/bindings/binding-state-v4relay-binding** and **relay/bindings/binding-state-v4relay-bound** are not correct. PR1475248

MPLS

- The FPC might be stuck in the **Ready** state after making a change in the configuration that removes RSVP and triggers FPC restart. PR1359087
- On the MPC10E or MPC11E line card, the LDP and BFD sessions are dropped when the **fast-lookup-filter** has a default term with only accept as action and it is attached to the lo0 interface. PR1474204
- The root XML tag in the output is changed from **rsvp-pop-and-fwd-info** to **rsvp-pop-and-fwd-information** to be consistent with the XML tag convention. PR1365940
- Traffic is silently discarded after the LSP protection link on the third-party transit router goes down. PR1439251
- On the MPC10E line card, the P2MP LSP traceroute is not working. PR1440636
- The traffic might be silently discarded after the LACP times out. PR1452866
- P2MP LSP might flap after the VT interface in the MVPN routing instance is reconfigured. PR1454987
- The rpd core files are generated with SNMP polling. PR1457681
- All LDP adjacencies flap after changing LDP preference. PR1459301
- The previously configured credibility preference is not considered by CSPF even though the configuration has been deleted or changed to prefer another protocol in the traffic engineering database. PR1460283
- MPLS trace route does not trace the SRUDP tunnel ingress router. PR1460516
- The process rpdtmd might crash while SNMP polls the statistics of the lpd interface. PR1465729
- The device might use the locally computed path for the PCE-controlled LSPs after the link or node fails. PR1465902
- The fast reroute detour next-hop down event might cause the primary LSP to go in the **Down** state in a particular scenario. PR1469567
- The p2mp traceroute fails with an aggregated Ethernet bundle over AFT. PR1470815
- The rpd process might crash during shutdown. PR1471191
- The rpd crash might be seen after some commit operations, which might affect the RSVP ingress routes. PR1471281
- The following error messages continuously floods the backup Routing Engine: **(JTASK_IO_CONNECT_FAILED: RPDTM./var/run/rpdtmd_control: Connecting to 128.0,255.255,255.255,0.0.0.0,0.0.0.0, failed: No such file or directory).** PR1473846
- RSVP LSPs might not come up in the scaled network with a very high number of LSPs if NSR is used on the transit router. PR1476773
- Kernel crashes and device might restart. PR1478806
- The rpd process crashes on the backup Routing Engine when LDP tries to create LDP p2mp tunnel upon receiving corrupted data from the master Routing Engine. PR1479249

Network Management and Monitoring

- The SNMP cold start trap might be seen after the Routing Engine switchover. PR1461839

Platform and Infrastructure

- The jcrypto syslog help package and events are not packaged even when the error message is compiled. PR1290089
- The time convergence for the MVPN fast upstream failover might be more than 50 minutes. PR1478981
- With chained composite next-hop enabled, the MPLS CoS rewrite does not work for IPv6 PE device traffic. PR1436872
- In an EVPN-VXLAN scenario, sometimes the host-generated packets get dropped when hitting the reject route in the Packet Forwarding Engine. PR1451559
- The MPC might drop packets after enabling the firewall fast lookup filter. PR1454257
- Multicast traffic loss occurs in a rare case in a seamless MPLS with MVPN configuration. PR1456905
- Port mirroring does not occur with VPLS. PR1458856
- DDoS-protection does not stop logging when the remote tracing is enabled. PR1459605
- Traceroute initiated from the PE device does not show the tunnel endpoint hop in the output. PR1461441
- CLI configuration flag **version-03** must be optional. PR1462186
- On the MX204 platform, Packet Forwarding Engine errors might occur when the incoming GRE tunnel fragments get sampled and undergo inline reassembly. PR1463718
- Not able to view the snapshots of the backup Routing Engine. PR1464394
- MX80 EVPN-VXLAN RT5 does not work properly, and **ip-prefix-routes** are not reachable. PR1466602
- On the MX150 devices, the default subscriber management license does not include the Layer 2 TP. PR1467368
- On the MX Series Virtual Chassis, the Layer 2 traffic sent from one member to another member is corrupted. PR1467764
- The JNH memory leaks after the CFM session flap for the LSI and VT interfaces. PR1468663

Routing Policy and Firewall Filters

- Routes resolution might be inconsistent if any route resolves over the multipath route. PR1453439

Routing Protocols

- The CPU utilization on rpd spins at 100 percent once the same external BGP route is learned on different VRF tables. PR1442902
- If the same neighbor is configured under different RIP groups, the commit check fails to capture this invalid configuration and commit is done successfully. However, the rpd process crashes. PR1485009
- The rpd crash might be seen after configuring OSPF **nssa area-range** and summaries. PR1444728

- The BGP routes might fail to be installed in a routing instance if the **from next-hop** policy match condition is used in the VRF import policy. PR1449458
- TI-LFA backup path for the adj-sids is broken in OSPF, where the shortest path to the node opposite the adj-sid is not the one-hop path over the interface indicated by the adj-sid. PR1452118
- The SSH login might fail if a user account exists in both the local database and RADIUS/TACACS+. PR1454177
- The rpd scheduler slip for BGP GR might be up to 120 seconds after the peer goes down. PR1454198
- MoFRR with MLDP inband signaling is not working. PR1454199
- The rpd memory might leak in certain MSDP scenario. PR1454244
- The rpd might crash continuously due to memory corruption in the IS-IS setup. PR1455432
- Packet drop and CPU spike on the Routing Engine might be seen in certain conditions if **labeled-unicast protection** is enabled for a CsC-VRF peer. PR1456260
- The topology-independent loop-free alternate might be unable to install backup path in the routing table in a specific case. PR1458791
- The rpd memory leak might be observed on the backup Routing Engine due to BGP flap. PR1459384
- The other querier present interval timer cannot be changed in a IGMP or MLD snooping scenario. PR1461590
- The rpd scheduler slips might be seen on the RPKI route validation enabled BGP peering router in a scaled setup. PR1461602
- Need to install all possible next hops for the OSPF network LSAs. PR1463535
- The IS-IS IPv6 multitopology routes might flap every time when there is an unrelated commit under the protocol statement. PR1463650
- The rpd might crash if both the BGP add-path and BGP multipath are enabled. PR1463673
- The rpd might crash if the IPv4 routes are programmed with the IPv6 next hop via JET APIs. PR1465190
- The BGP peers might flap if the **hold-time** parameter is set as small. PR1466709
- The configured BGP damping policy might not take effect after BGP is disabled and then enabled followed by commit. PR1466734
- BGP multipath does not work for MT on cRPD. PR1467091
- The rpd might crash after configuring **independent-domain** under the master routing instance. PR1469317
- The mcsnoopd might crash when the STP moves the mrouter port to the **Blocked** state. PR1470183
- The BFD client session might flap when removing the BFD configuration from the peer end (from other vendor) of the BFD session. PR1470603
- The rpd might crash when both the **instance-import** and **instance-export** policies contain the **as-path-prepend** action. PR1471968

- The rpd process might crash with the BGP multipath and damping configured. PR1472671
- Removal of the cluster from the BGP group might cause prolonged convergence times. PR1473351
- SFTP does not connect properly and the following error message is seen**: Received message too long.** PR1475255
- The rpd process might crash with BGP multipath and route withdrawal occasionally. PR1481589
- Removal of the BGP and rib-sharding configuration might cause the routing protocols to become unresponsive. PR1485720
- High CPU utilization might be observed when the outgoing BGP updates are sent slowly. PR1487691

Services Applications

- The jl2tpd process might crash during the restart procedure. PR1461335
- The calling station gets truncated after 64 bytes. PR1462689
- On an MX Series router, L2tp LTS fails to forward the **agentCircuitId** and **agentRemoteId** AVP toward the LNS. PR1472775
- Phase 1 SA migrates to a new remote IP because of the **source-address translation** for the static NAT tunnel. PR1477181

Subscriber Access Management

- The authd crashes on the backup Routing Engine during execution of the slax script that runs the **< get-jsrc-counters>** RPC call. PR1458185
- DHCPv6 subscribers might be stuck in a state after the authd process crashes. PR1460578
- A problem arises with **linked-pool-aggregation** after attempting to delete a pool in the middle of the chain. PR1465253
- The volume statistics attributes are missing in the accounting-stop for the Configuration Activated Services and CLI Activated Services. PR1470434
- The sub interfaces might be missing in the NAS port ID. PR1472045
- The authd process might crash after the ISSU setup from Junos OS Release 18.3 and earlier to Junos OS Release 18.4 and later. PR1473159
- Some address-relevant fields are missing when executing the **test aaa ppp** command. PR1474180
- The CoA request might not be processed if it includes the **proxy-state** attribute. PR1479697
- The **mac-address** CLI option is hidden under the **access profile radius options calling-station-id-format** statement. PR1480119

User Interface and Configuration

- On an MX Series device, a J-Web page might not get redirected to login once the session expires with an idle timeout. PR1459888

VPNs

- The P1 configuration delete message is not sent on loading baseline configuration if there has been a prior change in VPN configuration. PR1432434
- The rpd process might crash due to memory leak in MVPN RPF Src PE block. PR1460625
- The Layer 2 circuit displays MM status, which might cause traffic loss. PR1462583
- The Layer 2 circuit connections might become stuck in the **OL** state after changing the Layer 2 circuit community and flapping the primary LSP path. PR1464194
- The rpd might crash when **link-protection** is added or deleted from LSP for the MVPN ingress replication selective provider tunnel. PR1469028

SEE ALSO

What's New | **95**

What's Changed | **125**

Known Limitations | **133**

Open Issues | **138**

Documentation Updates | **211**

Migration, Upgrade, and Downgrade Instructions | **212**

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 20.1R3 documentation for the MX Series.

Dynamic Host Configuration Protocol (DHCP)

- **Introducing DHCP User Guide**—Starting in Junos OS Release 20.1R1, we are introducing the DHCP User Guide for Junos OS routing, switching, and security platforms. This guide provides basic configuration details for your Junos OS device as DHCP Server, DHCP client, and DHCP relay agent.

  [See DHCP User Guide.]

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the 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 20.1R2

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-x86-32-20.1R2.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-20.1R2.9-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-x86-32-20.1R2.x-limited.tgz
```

- For 64-bit Routing Engine version:

  `user@host>` **request system software add no-validate reboot *source*/junos-install-mx-x86-64-20.1R2.9-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 20.1R2, 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 20.1R2 **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-20.1R2.9-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-20.1R2.9-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 20.1R2 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Upgrading a Router with Redundant Routing Engines

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

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

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

Downgrading from Release 20.1R2

To downgrade from Release 20.1R2 to another supported release, follow the procedure for upgrading, but replace the 20.1R2 jinstall package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the Installation and Upgrade Guide.

SEE ALSO

Junos OS Release Notes for NFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

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

NOTE: For information about NFX product compatibility, see NFX Product Compatibility.

What's New in Release 20.1R2

There are no new features or enhancements to existing features for NFX Series devices in Junos OS Release 20.1R2.

What's New in Release 20.1R1

Application Security

- **AppQoE support for granular APBR rules (NFX Seris, SRX Series, and vSRX)**—Starting in Junos OS Release 20.1R1, AppQoE utilizes the granular rule matching functionality of advanced policy-based routing (APBR) for better quality of experience (QoE) for the application traffic.

  In Junos OS Release 18.2R1, APBR supported configuring policies by defining source addresses, destination addresses, and applications as match conditions. After a successful match, the configured APBR profile is applied as an application services for the session. In this release, AppQoE leverages the APBR enhancement and selects the best possible link for the application traffic as sent by APBR to meet the performance requirements specified in SLA.

  [See Application Quality of Experience.]

- **Default mechanism to forward the traffic through APBR rule (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS 20.1R1, you can configure a APBR rule by specifying the dynamic application match criteria with any keyword. This provides a default mechanism to forward the traffic to a specific next-hop device or to a destination if the traffic matches any dynamic application.

  [See Advanced Policy-Based Routing.]

- **Custom application enhancements (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.1R1, we've enhanced the custom applications signature functionality by providing a new set of applications and contexts.

  Application identification allows you to create custom application signatures to detect applications specific to your network environment. You can create custom application signatures for applications based on ICMP, IP protocol, IP address, and Layer 7 or TCP/UDP stream. While configuring the custom application signatures, you must specify the context values that the device can use to match the patterns in the application traffic.

  Custom application signature contexts are part of application signature package. You must download and install the latest application signature package version 3248 or later to use new contexts for custom application signatures.

  [See Custom Application Signatures for Application Identification.]

Interfaces

- **Single-leg and unidirectional cross-connect**— Starting in Junos OS Release 20.1R1, NFX Series devices support single-leg cross-connect and unidirectional cross-connect features.

  Single-leg cross-connect feature allows configuration of single entry in the cross-connect. The entry can be either VNF interface or a virtual interface. You can configure the other entry in the cross-connect at any later point of time.

  Unidirectional cross-connect feature allows the traffic to be forwarded conditionally or unconditionally in a single direction. Traffic flow in the opposite (other) direction follows the MAC-based forwarding rule.

  [See How to Configure NFX150, How to Configure NFX250, and How to Configure NFX350.]

Virtualized Network Functions (VNFs)

- **Virtual router reflector (VRR) virtualized network function (VNF) in enhanced orchestration (EO) mode**—Starting in Junos OS Release 20.1R1, you can instantiate the VRR VNF in EO mode by using the JDM CLI configuration and without using the XML descriptor file. EO mode uses Open vSwitch (OVS) as an NFV backplane for bridging the interfaces.

  [See Managing Virtual Network Functions Using JDM.]

SEE ALSO

What's Changed | 222

Known Limitations | 223

Open Issues | 223

Resolved Issues | 224

What's Changed

IN THIS SECTION

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

What's Changed in Release 20.1R3

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in Junos OS Release 20.1R3 for NFX Series devices.

What's Changed in Release 20.1R2

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in Junos OS Release 20.1R2 for NFX Series devices.

What's Changed in Release 20.1R1

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in Junos OS Release 20.1R1 for NFX Series devices.

SEE ALSO

Known Limitations

There are no known limitations for NFX Series devices in Junos OS Release 20.1R3.

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

SEE ALSO

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Release 20.1R3 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.

Platform and Infrastructure

- On NFX250 devices, vector packet processing (VPP) is not loaded in dual CPE, and at times in single CPE. PR1461238

Virtual Network Functions (VNFs)

- On NFX Series devices, while configuring **vmhost vlans** using **vlan-id-list**, the system allows duplicate VLAN IDs in the VLAN ID list. PR1438907
- On NFX Series devices, analyzers can be configured on VNF interfaces with output port as other VNF interfaces. All the packets entering or exiting can be mirrored on to the designated analyzer port. It is observed that after a system reboot, this functionality stops working and no packets are mirrored on the output analyzer port. PR1480290

SEE ALSO

What's New | **220**

What's Changed | **222**

Known Limitations | **223**

Resolved Issues | **224**

Documentation Updates | **228**

Migration, Upgrade, and Downgrade Instructions | **228**

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS Release 20.1R3 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: 20.1R3

Platform and Infrastructure

- False positive TSensor errors are reported on vjunos0. PR1508580
- When you upgrade the NFX150 devices from Junos OS Release 19.4 to Junos OS Release 20.2, the upgrade fails and an error message is displayed: **"/usr/sbin/boot_mgmt_fsm: line 40: echo: write error: No space left on device issue"**. PR1532334
- The control link might be broken when there is excessive traffic load on the control link in vSRX cluster deployment. PR1524243
- On NFX150 devices, when J-Flow v5 is configured and the J-Flow v5 server is reachable through an IPsec 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
- The l2cpd core files might be seen on reboot. PR1561235

Routing Protocols

- On NFX-series and MX150 devices 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

Resolved Issues: 20.1R2

Application Security

- AppQoE is sending active probe packets for the deleted **active-probe-params**. PR1492208

Interfaces

- On NFX250 NextGen devcies, the **monitor interface traffic** command might not display the output pps for sxe and physical interfaces. PR1464376
- On NFX350 devices, the **show interfaces | no-more** command output freezes for 20 seconds after displaying the **dl0** interface information. PR1502626

Platform and Infrastructure

- The device reads the board ID from EEPROM directly using I2C upon power cycle. PR1529667
- On NFX150 devices, ZTP over LTE configuration commit fails for **operation=create** in an XML operations configuration. PR1511306
- After you upgrade the JDM image from Junos OS Release 15.1X53-D497.1 to Junos OS Release 18.4R3-S2, tunnels are down in the gateway router. PR1507165

- On NFX150 devices, MAC aging does not work. You must remove aged MAC entries from the CLI. PR1502700
- The request vmhost power-off command reboots the NFX250 NextGen device instead of powering off the device. PR1493062
- After initiation of zeroization, the NFX250 device is going into a reboot loop. PR1491479
- Core files on NFX250 while adding the second LAN subnet. PR1490077
- Potential security vulnerabilities in Intel firmware that is used in the NFX150 network services platform may allow escalation of privilege, denial of service, or information disclosure. Intel has released firmware updates to mitigate these potential vulnerabilities. PR1480976
- The Juniper Device Manager (JDM) container, used by the disaggregated Junos OS architecture on Juniper Networks NFX350 devices, stores password hashes in the world-readable **file /etc/passwd**. This is not a security best practice as it can allow an attacker with access to the local file system the ability to decrypt password hashes stored on the system by brute force. PR1462556

Resolved Issues: 20.1R1

High Availability

- On an NFX150 high availability chassis cluster, the host logs updated in the system log messages might not show the correct time stamp. PR1394778

Interfaces

- On NFX150 devices, no error is displayed when the commit fails after you configure **native-vlan-id** on an access VNF interface. PR1438854
- On NFX Series devices, ping is not working between the cross-connected interfaces with interface deny-forwarding configuration. PR1442173
- On NFX Series devices, the static MAC address is replaced by a random MAC address. PR1458554
- When traffic goes through vSRX3.0 platforms, core-dump files are generated and traffic is dropped. This issue might result in the Packet Forwarding Engine being inactive and all interfaces being down. PR1465132
- On NFX150 devices, GRE tunnel interface (gr-1/0/0) might not appear if the **clear-dont-fragment-bit** option is configured for the GRE interface. PR1472029
- On NFX350 devices, if you delete and add SXE interfaces, the SXE interface moves to Spanning Tree Protocol blocking (STP BLK) state, and the traffic drops on that interface. PR1475854

Mapping of Address and Port with Encapsulation (MAP-E)

- On NFX Series devices, IP identification (IP ID) is not changed after MAP-E NAT44 is performed on fragment packets when the packets reach the customer edge (CE) device.

  PR1478037

Platform and Infrastructure

- LTE package related files are lost after image upgrade from Junos OS Release D497.1 to Junos OS Release 18.4R3.3 on NFX250 devices. PR1493711
- On NFX Series devices, if there are any conditional groups, the l2cpd process might crash and generate a core dump when interfaces are flapping and the LLDP neighbors are available. It might cause the dot1x process to fail and all the ports have a short interruption at the time of process crash. PR1431355
- Half-duplex configuration on 1-Gigabit Ethernet ports is not supported when auto negotiation is disabled. PR1453911
- On NFX350 devices, if you execute the **show vmhost mode** command multiple times, JDM might crash and cause the **show vmhost mode** commands to stop working. PR1474220
- After a power outage, JDMD is not responsive because the **/etc/hosts** file is corrupted on NFX250 devices. PR1477151

Routing Protocols

- On NFX Series devices, changing the **other querier present interval** timer is not working on IGMP or the MLD snooping device in the existing bridge domain (BD) or listener domain (LD). PR1461590

Virtualized Network Functions (VNFs)

- On NFX150 and NFX250 NextGen devices, when two flowd interfaces are mapped to the same physical interface and if you delete the interface mapping to VF0, the traffic flow is disrupted. Even though the mapping is moved to VF0, the MAC address is not cleared in VF1, which disrupts the traffic. PR1448595
- On NFX150 devices, when you need to change the vmhost mappings of a particular NIC or NICs, you must delete the existing vmhost mapping and commit the configuration. Now you can configure the new mappings for the respective NICs. You cannot change the NIC vmhost mappings in the same commit to delete and add a new mapping to the heth NICs. PR1459885
- NFX250 devices do not allow *jdm* (case-insensitive) as a VNF name. You can use *jdm* as a part of the name. For example, *jdm123*, *abcJDM*, *abcJDM123* are valid VNF names, whereas, *jdm*, *JDM*, *Jdm*, *JDm* are not valid VNF names. PR1463963

SEE ALSO

What's New | **220**

What's Changed | **222**

Documentation Updates

There are no errata or changes in Junos OS Release 20.1R3 documentation for NFX Series.

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

NOTE: For information about NFX product compatibility, see NFX Product Compatibility.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

Basic Procedure for Upgrading to Release 20.1

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 20.1R3:

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

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

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

3. Select the **Software** tab.

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

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

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

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

9. Copy the software to the device or to your internal software distribution site.

10. Install the new package on the device.

SEE ALSO

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 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 PTX Series.

What's New in 20.1R3

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 20.1R3.

What's New in 20.1R2

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 20.1R2.

What's New in 20.1R1

Interfaces and Chassis

- **Handling thermal health events (PTX5000)**—Starting in Junos OS Release 20.1R1, on PTX5000 routers, you can enable a thermal health check and configure an action (such as auto shutdown and alarm) to be taken when a thermal health event such as power leakage is detected. You can also configure the power supply module (PSM) watchdog to shut down the PSM output power in case a thermal health event causes Junos to go down.

  > **NOTE:** The PSM watchdog feature works only if all the online PSMs in the router support this feature.

  [See Handling Thermal Health Events Using Thermal Health Check and PSM Watchdog]

- **Support for new show | display set CLI commands (ACX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the following new **show** commands have been introduced:

  - **show | display set explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the top level of the hierarchy.
  - **show | display set relative explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the current hierarchy level.

  [See show | display set and show | display set relative.]

Junos OS XML API and Scripting

- **The jcs:load-configuration template supports loading the rescue configuration (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **jcs:load-configuration** template supports the **rescue** parameter to load and commit the rescue configuration on a device. SLAX and XSLT scripts can call the **jcs:load-configuration** template with the **rescue** parameter set to **"rescue"** to replace the active configuration with the rescue configuration.

  [See Changing the Configuration Using SLAX and XSLT Scripts and jcs:load-configuration Template.]

Junos Telemetry Interface

- **IS-IS adjacency and LSDB event streaming support on JTI (MX960, PTX1000, and PTX10000)**—Junos OS Release 20.1R1 provides IS-IS adjacency and link-state database (LSDB) statistics using Junos telemetry interface (JTI) and remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. ON_CHANGE statistics are sent to an outside collector.

  The following resource paths are supported:

- **/network-instances/network-instance/protocols/protocol/isis/interfaces/interfaces/circuit-counters/state/** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/interfaces/interface/levels/level/packet-counters/** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/system-level-counters/state/** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/interfaces/interfaces/levels/level/adjacencies/adjacency/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-ipv4-reachability/prefixes/prefix/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-ipv4-reachability/prefixes/prefix/subtlvs/subtlv/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv6-reachability/prefixes/prefix/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv6-reachability/prefixes/prefix/subtlvs/subtlv/** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-is-reachability/neighbors/neighbors/subTLVs/subTLVs/adjacency-sid/sid/state/** (ON-CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/extended-is-reachability/neighbors/neighbors/subTLVs/subTLVs/lan-adjacency-sid/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv4-interfaces-addresses/state/** (ON_CHANGE))
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv4-srlg/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv4-te-router-id/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/ipv6-interfaces-addresses/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/router-capabilities/router-capability/subtlvs/subtlv/segment-routing-capability/state/** (ON_CHANGE)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/state** (stream)

- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/area-address/state/address** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/nlpid/state/nlpid** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/lsp-buffer-size/state/size** (stream)
- **/network-instances/network-instance/protocols/protocol/isis/levels/level/link-state-database/lsp/tlvs/tlv/hostname/state/hname** (stream)

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

- **Platform, interface, and alarm sensor ON_CHANGE support on JTI (MX960, MX2020, PTX1000, PTX5000)**—Junos OS Release 20.1R1 supports platform, interface, and alarm statistics using Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services. You can use this feature to send ON_CHANGE statistics for a device to an outside collector.

  This feature supports the OpenConfig models:

  - **openconfig-platform.yang**: oc-ext:openconfig-version 0.12.1
  - **openconfig-interfaces.yang**: oc-ext:openconfig-version 2.4.1
  - **openconfig-alarms.yang**: oc-ext:openconfig-version 0.3.1

  Use the following resource paths in a gNMI subscription:

  - **/components/component** (for each installed FRU)
  - **/interfaces/interface/state/**
  - **/interfaces/interface/subinterfaces/subinterface/state/**
  - **/alarms/alarm/**

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

- **gRPC Dial-Out support on JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Junos OS Release 20.1R1 provides remote procedure call (gRPC) dial-out support for telemetry. In this method, the target device (server) initiates a gRPC session with the collector (client) and, when the session is established, streams the telemetry data that is specified by the sensor-group subscription to the collector. This is in contrast to the gRPC network management interface (gNMI) dial-in method, in which the collector initiates a connection to the target device.

  gRPC dial-out provides several benefits as compared to gRPC dial-in, including simplifying access to the target advice and reducing the exposure of target devices to threats outside of their topology.

  To enable export of statistics, include the **export-profile** and **sensor** statements at the [**edit services analytics**] hierarchy level. The export profile must include the reporting rate, the transport service (for example, gRPC), and the format (for example, gbp-gnmi). The sensor configuration should include the

name of the collector (the server's name), the name of the export profile, and the resource path. An example of a resource path is **/interfaces/interface[name='fxp0'**.

[See Using gRPC Dial-Out for Secure Telemetry Collection.]

- **gRPC version v1.18.0 with JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Junos OS Release 20.1R1 includes support for remote procedure call (gRPC) services version v1.18.0 with Junos telemetry interface (JTI). This version includes important enhancements for gRPC. In earlier releases, JTI is supported with gRPC version v1.3.0.

  Use gRPC in combination with JTI to stream statistics at configurable intervals from a device to an outside collector.

  [See gRPC Services for Junos Telemetry Interface.]

- **LLDP statistics, notifications, and configuration model for suppress-tlv-advertisement support on JTI (MX240, MX480, MX960, MX10003, PTX10008, PTX10016)**—Junos OS Release 20.1R1 provides remote procedure call (gRPC) streaming services support for attribute leaf **suppress-tlv-advertisement** under the resource path **/lldp/state/suppress-tlv-advertisement**. The following TLVs are supported, which in turn support operational state, notifications, and configuration change support:

  - port-description
  - system-name
  - system-description
  - system-capabilities
  - management-address
  - port-id-type

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

- **SR-TE statistics for uncolored SR-TE policies streaming on JTI (MX Series, PTX Series)**—Junos OS Release 20.1R1 provides segment routing traffic engineering (SR-TE) per label-switched Path (LSP) route statistics using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. Using JTI and gRPC services, you can stream SR-TE telemetry statistics for uncolored SR-TE policies to an outside collector.

  Ingress statistics include statistics for all traffic steered by means of an SR-TE LSP. Transit statistics include statistics for traffic to the Binding-SID (BSID) of the SR-TE policy.

  To enable these statistics, include the **per-source per-segment-list** statement at the [**edit protocols source-packet-routing telemetry statistics**] hierarchy level.

  If you issue the **set protocols source-packet-routing telemetry statistics no-ingress** command, ingress sensors are not created.

  If you issue the **set protocols source-packet-routing telemetry statistics no-transit** command, transit sensors are not created. Otherwise, if BSID is configured for a tunnel, transit statistics are created.

The following resource paths (sensors) are supported:

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

To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC.

Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface, source-packet-routing, and show spring-traffic-engineering lsp detail name *name*.]

Routing Protocols

- **Support for topology-independent loop-free alternate (TI-LFA) in IS-IS for IPv6-only networks (ACX Series, MX Series, and PTX Series)**— Starting with Junos OS Release 20.1R1, you can configure TI-LFA with segment routing in an IPv6-only network for the IS-IS protocol. TI-LFA provides MPLS fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. TI-LFA provides protection against link failure, and node failure.

  You can enable TI-LFA for IS-IS by configuring the **use-post-convergence-lfa** statement at the **[edit protocols isis backup-spf-options]** hierarchy level. You can enable the creation of post-convergence backup paths for a given IPv6 interface by configuring the **post-convergence-lfa** statement at the **[edit protocols isis interface *interface-name* level *level*]** hierarchy level. The **post-convergence-lfa** statement enables link-protection mode.

  You can enable **node-protection** mode for a given interface at the **[edit protocols isis interface *interface-name* level *level* post-convergence-lfa]** hierarchy level. However, you cannot configure fate-sharing protection for IPv6-only networks.

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

MPLS

- **Support for segment routing over RSVP forwarding adjacency (MX Series and PTX Series)**—Starting with Junos OS Release 20.1R1, we provide support for segment routing traffic to be carried over RSVP LSPs that are advertised as forwarding adjacencies in IS-IS. This feature is implemented in a network having LDP on the edge and RSVP in the core where you can easily replace LDP with IS-IS segment routing because it eliminates the need for MPLS signaling protocols such as LDP. This helps to remove a protocol from the network and results in network simplification.

  [See Understanding Segment Routing over RSVP Forwarding Adjacency in IS-IS.]

- **CoS-based forwarding and policy-based routing to steer selective traffic over an SR-TE path (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 20.1R1, you can use CoS-based forwarding (CBF) and policy-based routing (PBR, also known as filter-based forwarding or FBF) to steer service traffic using a particular segment routing-traffic-engineered (SR-TE) path. This feature is supported only on non-colored segment routing LSPs that have the next hop configured as a first hop label or an IP address.

With CBF and PBR, you can:

- Choose an SR-TE path on the basis of service.
- Choose the supporting services to resolve over the selected SR-TE path.

[See Example: Configuring CoS-Based Forwarding and Policy-Based Routing For SR-TE LSPs.]

Network Management and Monitoring

- **Remote port mirroring to an IP address (GRE encapsulation) (PTX Series)**—You use port mirroring to send traffic to applications that analyze traffic to monitor compliance, enforce policies, detect intrusions, and so on. Starting in Junos OS Release 20.1R1, you can configure remote port mirroring to send sampled packets to a remote IP address, with the packets encapsulated in a GRE header.

  - Configure remote port mirroring to send sampled packets to a remote IP address, with the packets encapsulated in an IPv4 GRE header:

    set forwarding-options port-mirroring instance *instance-name* output ip-source-address *address* ip-destination-address *address*

  - (Optional) Configure a static traffic-class value that represents the 8-bit differentiated services (DS) field in the IPv4 header of a GRE tunnel. You can program 6 of the 8 bits, so the value that you can configure under DSCP can be 0-63 (2^0 to 2^6).

    set forwarding-options port-mirroring instance *instance-name* output dscp *numeric-dscp-value*

  - (Optional) Configure a policer to police the mirrored traffic that is going out of that interface:

    set forwarding-options port-mirroring instance *instance-name* output policer *policer-name*

  - (Optional) Configure the forwarding of packets to a queue defined by a forwarding class:

    set forwarding-options port-mirroring instance *instance-name* output forwarding-class *forwarding-class-name*

  [See instance (Port Mirroring) and traffic-class (Tunnels).]

- **On-box monitoring support on the control plane (MX Series and PTX Series)**—Starting in Junos OS Release 20.1R1, you can configure on-box monitoring to monitor anomalies with respect to the memory utilization of Junos OS applications and the overall system in the control plane of MX Series and PTX Series routers.

  You can use on-box monitoring to monitor system-level memory and process-level memory to detect possible leaks. When the system is running low on memory, the process heuristic shares the prediction and you can configure the action to be taken when leaks are identified.

  See memory (system)

System Management

- **Restrict option under NTP configuration is now visible (ACX Series, QFX Series, MX Series, PTX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **noquery** command under the **restrict** hierarchy is now available and can be configured with a mask address. The **noquery** command is used to restrict ntpq and ntpdc queries coming from hosts and subnets.

  [See Configuring NTP Access Restrictions for a Specific Address.]

SEE ALSO

What's Changed

IN THIS SECTION

See what changed in this release for PTX Series routers.

What's Changed in 20.1R3 Release

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.]

- **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).]

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.]

Routing Protocols

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the lsdist.0 and lsdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, we added multiple secondary loopback addresses in the traffic engineering database to the lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router ID.

System Management

- **Support for exclude option under file archive (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—The **exclude** option is added under the command **file archive** that specifies the file pattern to exclude. This option helps to exclude files that delay compression or files that do not require compression.

  [See file archive.]

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.]

What's Changed in 20.1R2 Release

Class of Service (CoS)

- We've corrected the output of the "show class-of-service interface | display xml" command. Output of the following sort:
  <container>
  <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container> will now appear correctly as: <container> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3></container> <container> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container>.

General Routing

- **Displaying accurate aggregate drop statistics (MX Series)**—You can view the accurate aggregate drop statistics when a packet drop is seen on an aggregated Ethernet interface by using the **show interfaces extensive** command. In earlier releases, the **show interfaces extensive** command did not display accurate aggregate drop statistics. Only the individual aggregate child interface displayed accurate drop statistics.
- **Trigger alarms when a PTX10008 or PTX10016 router has a mix of AC and DC power supplies**—If you install a mix of AC and DC power supply units (PSUs), Junos OS raises an alarm to indicate that there is a mix of AC and DC power supplies in the router. To fix this alarm, you need to ensure that you install the same type of power supplies.

  [**See** Understanding Chassis Alarms.**]**
- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**—Starting in this release, we've renamed the **arp-snoop** packet type option in the **edit system ddos-protection protocols arp** protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial-of-service (DDoS) protection policer parameters for ARP traffic.

  [**See** protocols (DDoS)**.**]

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—You can set the verbosity of the trace log to only show error messages using the error option at the **edit system services extension-service traceoptions level** hierarchy.

  [**See** traceoptions (Services).]

MPLS

- **Change in auto bandwidth adjustment (PTX5000)**—If auto bandwidth adjustment fails because of bandwidth unavailable error, the router tries to bring up the LSP with the same bandwidth during the subsequent reoptimization. In earlier releases, when the auto bandwidth adjustment fails, the current bandwidth is reset to the bandwidth that was already active.

  [See rsvp-error-hold-time.]

What's Changed in 20.1R1 Release

There are no changes in behavior and syntax for EX Series in Junos OS Release 20.1R1.

SEE ALSO

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for PTX 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

- PTX Series platforms with the FPC-PTX-P1-A or FPC2-PTX-P1A line card might encounter a single event upset (SEU) event that can cause a linked-list corruption of the TQCHIP. The following syslog message is 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** Junos OS chassis management error handling does detect such condition, and raises an alarm and performs the **disable-pfe** action for the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, a restart of the FPC is needed. Soft errors are transient or non-recurring. FPCs experiencing such SEU events do not have any permanent damage. Contact your Juniper support representative if the issue is seen after a FPC restart. PR1254415
- In the specific case of semigraceful RCB reboot initiated by the internal shell command **vhclient init 0**, GRES takes longer than 3 minutes to complete as opposed to 21 seconds. As a workaround, the CLI command **request vmhost reboot** (graceful) and plugging out and plugging in the Routing Engine (ungraceful) do not exhibit this delay. PR1312065
- When a filter is attached in the outbound direction, GRE encapsulated headers are applied after the filter block in the egress direction. So in this case, it is possible that the filter is evaluated on an old header content (and not on the new GRE encapsulated header) and hence filter evaluation turns true and the new GRE encapsulated gets recirculated for another GRE encapsulation. This issue is difficult to fix as filter block evaluation happens before the new header is attached. PR1465837
- All multipath legs are not created in the presence of destination networks 0/0 and toggling between the BGP signal and IP-IP. PR1467022
- For scaled Macs, as per the current design, the learn rate is expected. PR1473334
- During reconfigurations/link events at the physical interface level, pe.ipw.misc_int.status:iq_disabled interrupts can be seen. These do not indicate impact to traffic. PR1476553
- PTX1000/PTX10000 platform count MPLS header also in packet length where as MX does not include it when acting in egress PE role. So we see difference in byte accounting in both platforms corresponding to the length of MPLS label stack received with the packet. PR1482408

MPLS

- Increasing ECMP from 64 to 128 may cause the ingress LSP setup rate to be lower due to increased number of next-hop changes for the IGP routes using a shortcut. PR1421976
- LDP session might drop during the FRR if the maxecmp is configured to 128 and LDP/IGP has more than 64 RSVP LSP next hops and LDP tunneling is configured on those next hops. PR1430361

- On all Junos OS platforms with distributed CSPF under SR-TE scenario, if you execute some operations such as deactivate or activate SR protocols, restart routing, and so on, rpd crash might be observed. PR1493721

SEE ALSO

What's New | 231

What's Changed | 238

Open Issues | 244

Resolved Issues | 247

Documentation Updates | 253

Migration, Upgrade, and Downgrade Instructions | 254

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Release 20.1R3 release for PTX 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 CFP2-DCO-T-WDM-1 is plugged in a PTX Series PIC, after FPC restarts, the carrier frequency offset TCA is raised even when TCA is not enabled. PR1301471
- On 30-Port MACSec linecard (LC1101-M - 30C / 30Q / 96X) of PTX10000 chassis, under certain circumstances, when **exclude-protocol lacp** configuration statement under [edit security macsec connectivity-association connectivity-association-name] hierarchy level is deleted or deactivated, the

LACP protocol "Mux State" shown under the output of the **show lacp interface** CLI command might remain as "attached" or "detached" and might not move to "distributing" state. PR1331412

- The QFX10000 platform drops the Aruba wireless access point (AP) heartbeat packets. As a result, the Aruba wireless AP cannot work. PR1352805
- Traffic loss is greater than 50 milliseconds (in order of 200 to 300 milliseconds) for IP routes pointing to unilist of composites with indirect next hops during a link down scenario. In this case, Packet Forwarding Engine do not local repair and will wait for the rpd to install the new next hops. PR1383965
- 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
- On routers and switches running Junos OS, with Link Aggregation Control Protocol (LACP) enabled, deactivating a remote aggregated Ethernet (AE) member link makes the local member link move to LACP detached state and cause traffic drops on that member link. The same scenario applied when a new member link is added where the other end of that link is not configured with LACP. PR1423707
- The em2 interface configuration causes FPC to crash during initialization and FPC does not come online. After deleting the em2 configuration and restarting the router, FPC comes online. PR1429212
- Memory leaks are expected in this release. PR1438358
- When users configure the best destination network with dyn-tunnel-attribute-policy and preference, the tunnel from the old destination network are not migrated. PR1462805
- During reconfigurations or link events at the physical interface level, pe.ipw.misc_int.status:iq_disabled interruption is seen. These do not indicate impact to traffic. PR1476553
- The Layer 2 VPN (L2VPN) on PTX Series with asynchronous-notification might keep flapping when the link is going up between PE and CE devices. After Layer 2 VPN flap, the interfaces which are set "asynchronous-notification" might show "- Inf dBm" laser output power even when the L2VPN status is up. PR1486181
- In strict priority scheduling mode, medium-high and medium-low are seen to be operating on the same priority. PR1490505
- Junos Telemetry GRPC multi sensors are not working as expected. PR1492282
- When **show interfaces aex extensive** statement is synchronized and SNMP polling queries are asynchronized on aggregated Ethernet interface in parallel, you might observe spikes in aggregated Ethernet interface framing errors counter in between correct values. PR1539537
- IS-IS over Layer 2 circuit might not come up if the encapsulation is TCC. PR1590387

Infrastructure

- On PTX Series platform the harmless log of "invalid SMART checksum" might be seen when performing software upgrade to specific releases (for example, Junos OS Releases 15.1F5-S3, 15.1F6-S1, 15.1F7, 15.1R4-S3, 15.1R5, 16.1R1, 16.1R2, 16.2R1). PR1222105
- Memory corruption of a binary from **/usr/bin/** or **/usr/sbin/** directory can occur if such binary is invoked when a recovery snapshot creation is in progress. The exact symptoms will be different depending on the exact binary and Junos OS version - some programs will show an error, and some programs will crash every time it is executed. Such memory corruption will be persistent until the affected Routing Engine is restarted. Refer to TSB17954 for more details. PR1563647

MPLS

- At high scale, LSP setup rate will be relatively slower in IP-in-IP networks. PR1457992

Routing Protocols

- With bidirectional forwarding detection (BFD) configured on an aggregated Ethernet interface, if you disable/enable the aggregated Ethernet interface, then that interface and the BFD session might not come up. PR1354409
- The **show dynamic-tunnels database** command does not show the current value of traffic statistics. It shows the cache value of traffic statistics, which might not be equal to the current value. PR1445705
- With NSR enabled, the current BGP design support 3000 BGP IPv6 peers or 8000 BGP IPv4 peers. When you are trying to bring up more than 3000 BGP IPv6 sessions or more than 8000 BGP IPv4 sessions, the rpd might crash. PR1461436

SEE ALSO

What's New | **231**

What's Changed | **238**

Known Limitations | **242**

Resolved Issues | **247**

Documentation Updates | **253**

Migration, Upgrade, and Downgrade Instructions | **254**

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in this release 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: 20.1R3

Forwarding and Sampling

- The l2ald process might crash because of the next hop issue in the EVPN-MPLS. PR1548124

General Routing

- Delay in disabling Packet Forwarding Engine might be seen on PTX Series devices with PECHIP equipped FPCs inserted. PR1481879
- The SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at the collector. PR1484322
- On PTX10008, FPC UKERN core file is not transferred to Routing Engine in scaled setup. PR1500418
- Error messages **t6e_dfe_tuning_state:et-6/0/0 - Failed to dfe tuning count 10** might be seen after links flap. PR1512919
- Packet drop might be seen with all commit events with 1G speed configured interface PR1524614
- The chassisd memory leak might cause traffic loss. PR1537194
- Observed 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** in PTX1000. PR1540064
- The Packet Forwarding Engine might crash in MPLS IPv6-tunneling scenario when the next hop changes. PR1540793
- Junos OS: PTX Series: Denial of Service in packet processing due to heavy route churn when J-Flow sampling is enabled (CVE-2021-0263). PR1546143

- Traffic might get discarded after swapping an FPC type 3 card with an FPC type 1 card in the same slot on a PTX3000 router. PR1547790
- The rpd crash might be seen when the BGP service route is resolved over the color-only SR-TE policy. PR1550736
- The interface filter with source-port 0 matches everything instead of port 0. PR1551305
- The LCMD process might consume memory until all the free memory available to VMHOST gets exhausted. PR1555386
- An enhancement to enable watchdog petting log on PTX10000 line cards. PR1561980
- TACACS traffic might be dropped. PR1578579
- **Failed to get pechip handle for chip 0** and **prds_encap_sample_flood_lpbk_desc_install: Egress NH descriptor install OK for Flabel 7808** errors are seen during bringup. PR1585594

Infrastructure

- Interface drop counters might display 0 during a race condition when voq statistics are also polled simultaneously. PR1537960
- Invalid statistics value might be observed when multiple mib2d/cosd requests for the same IFD arrive within 1 second. PR1541579
- The kernel crash with core file might be seen if churn happens for a flood composite next hop. PR1548545

Interfaces and Chassis

- Logs are not written in /var/log/messages on certain PTX Series routers running Junos Evolved. PR1551374

Multicast

- FPC might crash in a multicast scenario. PR1569957

Network Management and Monitoring

- The mib2d process crashes and generates a core file on backup Routing Engine. PR1557384

Platform and Infrastructure

- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. PR1552603
- FPC might crash in a scaled-firewall configuration. PR1586817

Routing Policy and Firewall Filters

- Generated route goes to the hidden state when the protect core command is enabled. PR1562867

Routing Protocols

- The BGP RPKI ROA withdrawal might lead to an unexpected BGP route flap. PR1483097
- The rpd might crash with BGP RPKI enabled in a race condition. PR1487486
- Traffic might be silently discarded when the BGP route gets deleted, which is part of multipath. PR1514966
- The rpd process generates the core file at **gp_rtarget_tsi_update,bgp_rtarget_flash_rt,bgp_rtarget_flash**. PR1541768
- BGP LU session flap might be seen with AIGP used scenario. PR1558102

User Interface and Configuration

- Any change in the nested groups might not be detected on commit and does not take effect. PR1484801

Resolved Issues: 20.1R2

General Routing

- On the PTX10008 and PTX5000 routers, the output of the **show filter index *number* counter** command shows value as zero at 28-02-HOSTBOUND_NDP_DISCARD_TERM. PR1420057
- The **show snmp mib walk jnxContentsDescr** command output does not display the fan controllers. PR1455640
- PHP device has Nnext-hop misprogramming for members of ECMP for SR label route used for reaching IPv6 destinations. PR1457230
- Interface statistics might not get updated for the local-loopback test. PR1458814
- PTX1000 and PTX10002 routers might drop packets after transient SIB or FPC voltage alarms. PR1460406
- On the PTX5000 routers, for the FPC3 line card, the optics-options syslog and link-down do not work as expected. PR1461404
- On the PTX10000 routers, FPCs might restart during runtime. PR1464119
- The router might become nonresponsive and bring the traffic down when the disk space becomes full. PR1470217
- A PTX5000 SIB3 might fail to come up in slot 0 and/or slot 8 when Routing Engine 1 is the primary Routing Engine. PR1471178
- Sampling process on FPC might crash when the MPLS traffic is sampled. PR1477445
- Multicast routes add or delete events might cause adjacency and LSPs to go down. PR1479789
- FPC might crash when dealing with invalid next hops. PR1484255

- BCM8238X SerDes firmware did not complete tuning; this might generate a false positive alarm. PR1491142
- BFD sessions start to flap when the firewall filter in loopback0 is changed. PR1491575
- PTX Series: Kernel routing table (KRT) queue stuck after J-Flow sampling of a malformed packet (CVE-2020-1679). PR1495788
- Outbound SSH connection flap or memory leak issue might be observed when pushing the configuration to the ephemeral database at a high rate. PR1497575
- The following error message is observed: **PFE_ERROR_FAIL_OPERATION: IFD et-1/0/8: RS credits failed to return: init=192 curr=193 chip=5**. PR1502716
- On the PTX10008 and PTX10016 routers, a few TCP-based application sessions might flap upon Routing Engine switchover or application sessions bouncing in the backup Routing Engine. PR1503169
- On the PTX3000 or PTX5000 router, unable to bring the ports up when plugging in the optic QSFP-100G-LR4-T2(740-061409). PR1511492
- The routes update might fail because of the HMC memory issue, and traffic impact might be seen. PR1515092
- Sampling with the rate limiter command enabled, crosses the sample rate 65,535. PR1525589

Interfaces and Chassis

- The following error message is observed: **Request failed: OID not increasing: ieee8021CfmStackServiceSelectorTyp**e. PR1517046
- EOAM IEEE802.3ah link discovery state is "Down" instead of "Active Send Local" after deactivating interfaces on routers. PR1532979

MPLS

- BGP session flap between two directly connected BGP peers because of the incorrect TCP-MSS in use. PR1493431
- The rpd process might crash in a rare condition in an SR-TE scenario. PR1493721
- The SNMP trap is sent with the incorrect OID jnxSpSvcSetZoneEntered. PR1517667

Routing Protocols

- The ppmd process crashes after configuring the S-BFD responder on the PTX Series routers with RE-DUO-2600. PR1477525
- The BGP multipath traffic might not fully load-balance for a while after adding a new path for load sharing. PR1482209
- The BGP route target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743
- The rpd might report 100 percent CPU usage with BGP route damping enabled. PR1514635

Resolved Issues: 20.1R1

Forwarding and Sampling

- The pfed might crash and be unable to come up on the PTX Series or TVP platforms. PR1452363

General Routing

- PTX Series interface stays down after maintenance. PR1412126
- Telemetry statistics might not account correctly when IS-IS sensors are enabled and the route next hops are ae interfaces. PR1413680
- LACP packet does not pass through Layer 2 circuit. PR1424553
- Interface does not come up after interface flapping and FPC reboot. PR1428307
- Reclassification policy applied on the route prefixes might not work on PTX Series platforms. PR1430028
- The l2cpd process might crash and generate a core file when interfaces are flapping. PR1431355
- The FPC might crash when a firewall filter is modified. PR1432116
- Unable to change DDoS protocol TTL values under PTX10000. PR1433259
- Upgrading fails due to communication failure between Junos VM and the host OS. PR1438219
- Packet loss might be seen if IPoIP or MPLS-over-UDP dynamic tunnels with ECMP are configured. PR1446132
- Changing the hostname triggers an on-change notification, not an adjacency on-change notification. Also, currently IS-IS is sending the hostname instead of the system ID in OC paths. PR1449837
- JNP10K-LC2101 FPC generates the "Voltage Tolerance Exceeded" major alarm for EACHIP 2V5 sensors. PR1451011
- The 100-Gbps interface might not come up after flapping on PTX5000. PR1453217
- Traffic might be dropped on PTX Series platforms. PR1459484
- Silent dropping of traffic upon interface flapping after DRD auto-recovery. PR1459698
- The "forwarding" option is missed in routing-instance type. PR1460181
- Hardware failure in CB2-PTX causes traffic interruption. PR1460992
- The **sample**, **syslog**, or **log** action in output firewall filters for packets of size less than 128 bytes might cause an ASIC wedge (all packet loss) on PTX Series platforms. PR1462634
- PIC might restart if the temperature of QSFP optics is overheated on PTX3000 or PTX5000. PR1462987
- An FPC might restart during runtime on PTX10000 or QFX10000 lines of devices. PR1464119
- Continuous MACsec-wedge-cleared logs might be seen and LACP flapping might happen with 100% line-rate traffic or near line rate traffic in the MACsec line card. PR1466481
- EBUF parity interrupt is not seen on PTX Series routers or the QFX10000 line of switches. PR1466532

- IPv6 traffic might get dropped in a Layer 3 VPN network. PR1466659
- Packet Forwarding Engine error logs (prds_packet_classify_notification: Failed to find fwd nh for flabel 48) might be reported when IGMP packets get sampled on the PTX5000 platform. PR1466995
- Optics measurements might not be streamed for interfaces of a PIC over JTI. PR1468435
- Incorrect counter value for **Arrival rate** and **Peak rate** for DDoS commands. PR1470385
- Traffic loops for pure Layer 2 packets coming over an EVPN tunnel with the destination MAC address matching the IRB MAC address. PR1470990
- The **input-vlan-map** or **output-vlan-map** might not work properly in a Layer 2 circuit local-switching scenario. PR1474876

Infrastructure

- The kernel crashes when removing a mounted USB storage device while a file is being copied to it. PR1425608
- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1 and later. PR1462986

Interfaces and Chassis

- After member interface flapping, the aggregated Ethernet remains down on the 5-port 100-Gigabit Ethernet DWDM CFP2-ACO PIC. PR1429279

Layer 2 Ethernet Services

- Member links state might be asychronized on a connection between the PE and CE devices in EVPN A/A scenario. PR1463791

MPLS

- Kernel crash and device restart might happen. PR1478806

Routing Protocols

- SSH login might fail if a user account exists in both local database and RADIUS or TACACS+. PR1454177
- The **other querier present interval** timer cannot be changed in an IGMP/MLD snooping scenario. PR1461590
- The rpd process might crash with BGP multipath and route withdrawal occasionally. PR1481589

SEE ALSO

What's New | **231**

What's Changed | **238**

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 20.1R1 documentation for the PTX Series.

Dynamic Host Configuration Protocol (DHCP)

- **Introducing DHCP User Guide**—Starting in Junos OS Release 20.1R1, we are introducing the DHCP User Guide for Junos OS routing, switching, and security platforms. This guide provides basic configuration details for your Junos OS device as DHCP Server, DHCP client, and DHCP relay agent.

  [See DHCP User Guide.]

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the 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 20.1

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 20.1R1:

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-x86-64-20.1R1.9.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-x86-64-20.1R1.9-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 20.1**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

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

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

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

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

Upgrading a Router with Redundant Routing Engines

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

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

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

SEE ALSO

What's New | 231

What's Changed | 238

Known Limitations | 242

Junos OS Release Notes for the QFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 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.

> **NOTE:** The following QFX Series platforms are supported in Release 20.1R3: QFX5100, QFX5110 (32Q and 48S), QFX5120, QFX5200, QFX5200-32CD, QFX5210, QFX10002, QFX10002-60C, QFX10008, and QFX10016.
>
> Junos on White Box runs on Accton Edgecore AS7816-64X switches in this release. The software is based on Junos OS running on QFX5210 switches, so release-note items that apply to QFX5210 switches also apply to Junos on White Box.

What's New in Release 20.1R3

There are no new features or enhancements to existing features for QFX Series in Junos OS Release 20.1R3.

What's New in Release 20.1R2

There are no new features or enhancements to existing features for QFX Series in Junos OS Release 20.1R2.

What's New in Release 20.1R1

EVPN

- **Routing traffic between a VXLAN and a Layer 3 logical interface (EX4650 and QFX5120)—**Starting in Junos OS Release 20.1R1, EX4650 and QFX5120 switches support the routing of traffic between a Virtual Extensible LAN (VXLAN) and a Layer 3 logical interface. This feature is enabled by default, so you do not need to take any action to enable it.

  > **NOTE:** By default, this feature is disabled on QFX5110 switches. To enable the feature on QFX5110 switches, you must perform the configuration described in Understanding How to Configure VXLANs and Layer 3 Logical Interfaces to Interoperate.

  (You can configure the Layer 3 logical interface using the **set interfaces *interface-name* unit *logical-unit-number* family inet address *ip-address/prefix-length*** or the **set interfaces *interface-name* unit *logical-unit-number* family inet6 address *ipv6-address/prefix-length*** command.)

High Availability (HA) and Resiliency

- **Inline keepalive packet support for BFD (QFX5110, QFX5120, QFX5200, and QFX5210)**—Starting in Junos OS Release 20.1R1, multihop BFD inline keepalive support enables scaling up to 10 inline BFD sessions with 150-millisecond support on both multihop BFD sessions as well as single-hop inline sessions. Multihop BFD session intervals can also be configured to less than 1-second granularity. This enables both faster detection of link failures and recovery. The switch will also send keepalive messages according to the configured interval.

  > **NOTE:** This feature only applies for IPv4 multihop BFD sessions and standalone BFD sessions. This feature is not supported for micro BFD implementation.

  [See Understanding Bidirectional Forwarding Detection (BFD).]

Interfaces and Chassis

- **Support for new show | display set CLI commands (ACX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the following new **show** commands have been introduced:
  - **show | display set explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the top level of the hierarchy.
  - **show | display set relative explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the current hierarchy level.

[See show | display set and show | display set relative.]

Junos OS XML, API, and Scripting

- **The jcs:load-configuration template supports loading the rescue configuration (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **jcs:load-configuration** template supports the **rescue** parameter to load and commit the rescue configuration on a device. SLAX and XSLT scripts can call the **jcs:load-configuration** template with the **rescue** parameter set to **"rescue"** to replace the active configuration with the rescue configuration.

  [See Changing the Configuration Using SLAX and XSLT Scripts and jcs:load-configuration Template.]

Junos Telemetry Interface

- **gRPC Dial-Out support on JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Junos OS Release 20.1R1 provides remote procedure call (gRPC) dial-out support for telemetry. In this method, the target device (server) initiates a gRPC session with the collector (client) and, when the session is established, streams the telemetry data that is specified by the sensor-group subscription to the collector. This is in contrast to the gRPC network management interface (gNMI) dial-in method, in which the collector initiates a connection to the target device.

  gRPC dial-out provides several benefits as compared to gRPC dial-in, including simplifying access to the target advice and reducing the exposure of target devices to threats outside of their topology.

  To enable export of statistics, include the **export-profile** and **sensor** statements at the [**edit services analytics**] hierarchy level. The export profile must include the reporting rate, the transport service (for example, gRPC), and the format (for example, gbp-gnmi). The sensor configuration should include the name of the collector (the server's name), the name of the export profile, and the resource path. An example of a resource path is **/interfaces/interface[name='fxp0'**.

  [See Using gRPC Dial-Out for Secure Telemetry Collection.]

- **gRPC version v1.18.0 with JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Junos OS Release 20.1R1 includes support for remote procedure call (gRPC) services version v1.18.0 with Junos telemetry interface (JTI). This version includes important enhancements for gRPC. In earlier releases, JTI is supported with gRPC version v1.3.0.

  Use gRPC in combination with JTI to stream statistics at configurable intervals from a device to an outside collector.

  [See gRPC Services for Junos Telemetry Interface.]

Multicast

- **PIM with IPv6 multicast traffic (EX4650 and QFX5120-48Y)**—Starting in Junos OS Release 20.1R1, EX4650 and QFX5120-48Y switches support Protocol Independent Multicast (PIM) with IPv6 multicast traffic as follows:
  - PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM sparse-dense mode (PIM-SDM)
  - PIM any-source multicast (PIM-ASM) and PIM source-specific multicast (PIM-SSM)

- Static, embedded, and anycast rendezvous points (RPs)

[See PIM Overview.]

Routing Policy and Firewall Filters

- **Support for flexible-match-mask match condition (EX4650 and QFX-Series)**—Starting with Junos OS Release 20.1R1, for EX4650, QFX5120-32C, and QFX5120-48Y switches, the **flexible-match-mask** match condition in firewall filters is supported for the **inet**, **inet6**, and **ethernet-switching** families. With this feature, you can configure a filter by specifying the length of the match (4 bytes maximum) starting from a Layer 2 or Layer 3 packet offset.

  [See Firewall Filter Flexible Match Conditions.]

Routing Protocols

- **Redistribution of IPv4 routes with IPv6 next hop into BGP (QFX Series)**—Starting in Release 20.1R1, devices running Junos OS can forward IPv4 traffic over an IPv6-only network, which generally cannot forward IPv4 traffic. As described in RFC 5549, IPv4 traffic is tunneled from CPE devices to IPv4-over-IPv6 gateways. These gateways are announced to CPE devices through anycast addresses. The gateway devices then create dynamic IPv4-over-IPv6 tunnels to remote CPE devices and advertise IPv4 aggregate routes to steer traffic. Route reflectors with programmable interfaces inject the tunnel information into the network. The route reflectors are connected through IBGP to gateway routers, which advertise the IPv4 addresses of host routes with IPv6 addresses as the next hop.

  To configure a dynamic IPv4-over-IPv6 tunnel, include the **dynamic-tunnels** statement at the **[edit routing-options]** hierarchy level.

  [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.]

Software Defined Networking

- **VMware NSX Data Center for vSphere 6.4.5 and 6.4.6 certification (QFX5100 Virtual Chassis)—**Starting with Junos OS Release 20.1R1, Juniper Networks certifies QFX5100 Virtual Chassis as a hardware Virtual Extensible LAN (VXLAN) gateway in an Open vSwitch Database (OVSDB) and VXLAN network with a VMware NSX Data Center for vSphere 6.4.5 or 6.4.6 controller.

  [See OVSDB-VXLAN User Guide for QFX Series Switches (VMware NSX).]

Storage and Fibre Channel

- **FIP snooping (EX4650-48Y and QFX5120-48Y)**—Starting in Junos OS Release 20.1R1, EX4650-48Y and QFX5120-48Y switches support Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) snooping. With FIP snooping enabled on these switches, you prevent unauthorized access and data transmission to a Fibre Channel (FC) network by permitting only those servers that have logged in to the FC network to access the network. You enable FIP snooping on FCoE VLANs when the switch is being used as an FCoE transit switch that connects FC initiators (servers) on the Ethernet network to FCoE forwarders at the FC storage area network (SAN) edge.

  [See Understanding FCoE Transit Switch Functionality and Understanding VN_Port to VN_Port FIP Snooping on an FCoE Transit Switch.]

System Management

- **Support for the Precision Time Protocol (PTP) AES67, SMPTE ST-2059-2, and AES67+SMPTE profiles (QFX10002)**—Starting in Junos OS Release 20.1R1, you can enable the AES67, SMPTE ST-2059-2, and AES67+SMPTE profiles to support video applications for capture (for example, cameras), video edit, and playback to be used in professional broadcast environments. The PTP standard allows multiple video sources to stay in synchronization across various equipment by providing time and frequency synchronization to all devices. These profile support PTP over IPv4 multicast and ordinary and boundary clocks.

  To configure the AES67, SMPTE ST-2059-2, and AES67+SMPTE profiles, enable one of the **aes67**, **smpte**, or **aes67-smpte** statements at the **[edit protocols ptp profile-type]** hierarchy level.

  [See Understanding the PTP Media Profiles.]

- **Restrict option under NTP configuration is now visible (ACX Series, QFX Series, MX Series, PTX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **noquery** command under the **restrict** hierarchy is now available and can be configured with a mask address. The **noquery** command is used to restrict ntpq and ntpdc queries coming from hosts and subnets.

  [See Configuring NTP Access Restrictions for a Specific Address.]

SEE ALSO

What's Changed

IN THIS SECTION

Learn about what changed in Junos OS main and maintenance releases for QFX Series.

What's Changed in 20.1R3

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.]

- **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).]

Layer 2 Ethernet Services

- **Link selection support for DHCP (QFX Series)**—We?ve introduced **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. Earlier to this release, the DHCP relay drops packets during the renewal DHCP process as the DHCP Server uses the leaf's address as a destination to acknowledge DHCP renewal message.

  [See relay-option-82.]

MPLS

- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the cspf-backoff-time statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

  [**See** cspf-backoff-time.]

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.]

Platform and Infrastructure

- 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)

Routing Protocols

- **Advertising 32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the lsdist.0 and lsdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, we added multiple secondary loopback addresses in the traffic engineering database to the lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router ID.

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.]

What's Changed in 20.1R2

Interfaces and Chassis

- **Autonegotiation status displayed correctly (QFX5120-48Y)**—In Junos OS Release 20.1R2, the **show interfaces *interface-name* <media> <extensive>** command displays the autonegotiation status only for the interface that supports autonegotiation. This is applicable when the switch operates at 1-Gbps speed.

  In the earlier Junos OS releases, incorrect autonegotiation status was displayed even when autonegotiation was disabled.

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—You can set the verbosity of the trace log to only show error messages using the error option at the edit system services extension-service traceoptions level hierarchy.

  [See traceoptions (Services)]

High Availability (HA) and Resiliency

- **IPv6 address in the prefix TIEs displayed correctly**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.
- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

Platform and Infrastructure

- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**— Starting in this release, we've renamed the **arp-snoop** packet type option in the **edit system ddos-protection protocols arp** protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial-of-service (DDoS) protection policer parameters for ARP traffic.

  See protocols (DDoS) (PTX Series and QFX Series) protocols (DDoS) (PTX Series and QFX Series).

- Priority-based flow control (PFC) support (QFX5120-32C)—Starting with Junos OS Release 20.1R2, QFX5120-32C switches support PFC using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic.

Routing Protocols

- **IGMP snooping in EVPN-VXLAN multihoming environments (QFX5110)**—In an EVPN-VXLAN multihoming environment on QFX5110 switches, you can now selectively enable IGMP snooping only on those VLANs that might have interested listeners. In earlier releases, you must enable IGMP snooping on all VLANs associated with any configured VXLANs because all the VXLANs share VXLAN tunnel endpoints (VTEPs) between the same multihoming peers and require the same settings. This is no longer a configuration limitation.

What's Changed in 20.1R1

Class of Service (CoS)

- We've corrected the output of the **show class-of-service interface | display xml** command. The output is of the following sort: <container> <leaf-1> data </leaf-1><leaf-2>data </leaf-2> <leaf-3> data</leaf-3> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container> will now appear correctly as <container> <leaf-1> data </leaf-1><leaf-2>data </leaf-2> <leaf-3> data</leaf-3></container> <container> <leaf-1> data </leaf-1> <leaf-2> data </leaf-2> <leaf-3> data </leaf-3> </container>.

Interfaces and Chassis

- **Commit error thrown when GRE interface and tunnel source interface are configured in different routing instances (QFX Series)**—In Junos OS Release 20.1R1, QFX Series switches do not support configuring the GRE interface and the underlying tunnel source interface in two different routing instances. If you try this configuration, it will result in a commit error with the following error message:

  error: GRE interface (gr-0/0/0.0) and its underlying tunnel source interface are in different routing-instances

  error: configuration check-out failed

  [See Understanding Generic Routing Encapsulation.]

- **Support for 100-Mbps speed using QFX-SFP-1GE-T on QFX5110-48S Switches**—Starting in Junos OS release 20.1R1, in addition to 1-Gbps, 10-Gbps, 40-Gbps, 100-Gbps speeds, now you can configure 100-Mbps speed using the **set interfaces *interface-name* speed 100M** command. By default, all 48 ports

on QFX5110-48S come up with 10-Gbps speed. With QFX-SFP-1GE-T connected, along with 1-Gbps speed, now you can also configure 100-Mpbs on QFX5110-48S switches.

[See Speed (Ethernet)].

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

Multicast

- **Multicast Layer 2 transit traffic statistics by multicast source and group (EX4600, EX4650, and the QFX5000 line of switches)**—Starting in Junos OS Release 20.1R1, EX4600, EX4650, and the QFX5000 line of switches provide statistics on the packet count for each multicast group and source when passing multicast transit traffic at Layer 2 with IGMP snooping. Run the **show multicast snooping route extensive** CLI command to see this count in the **Statistics: ... *n* packets** output field. The other statistics in that output field, **kBps** and **pps**, are not available (values displayed there are not valid statistics for multicast traffic at Layer 2). In earlier Junos OS releases, all three values in the **Statistics** output field for **kBps**, **pps**, and **packets** do not provide valid statistics for multicast traffic at Layer 2.

  [See show multicast snooping route.]

Network Management and Monitoring

- **entPhysicalTable fetched on QFX10002**—In Junos OS Release 20.1R1, the MIB data for entPhysicalTable will be fetched on a QFX10002-72Q or QFX10002-36Q switch.

  [See SNMP Explorer.]

Routing Protocol

- **Automatic installation of YANG-based CLI for RIFT protocol (MX Series, QFX Series, and vMX with 64-bit and x86-based servers)**—In Rift 1.2 Release, installation of the CLI for RIFT protocol occurs automatically along with the installation of the junos-rift package. In the pre-1.0 releases of the junos-rift package, the RIFT CLI had to be installed separately using **request system yang** command after installation of the junos-rift package.–

SEE ALSO

What's New | 258

Known Limitations | 270

Open Issues | 272

Resolved Issues | 277

Known Limitations

IN THIS SECTION

Learn about known limitations in Junos OS Release 20.1R3 for QFX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

Class of Service (CoS)

- On the QFX5100 line of switches, ISSU does not support Junos OS Release 20.1 and later. PR1479439
- Traffic might be dropped by the destination device. PR1568333

Infrastructure

- File system panic might occur after repeated power loss. PR1444941

Layer 2 Features

- On the QFX5000 line of switches, the following error message is reported in the log: **fpc0 Pools exhausted for Table:EGR_DVP_ATTRIBUTE_1**. PR1479826

Layer 2 Ethernet Services

- The LACP force-up and EVPN core isolation features are not supported together. PR1461581

Platform and Infrastructure

- Upgrade or downgrade from TVP to non-TVP is not supported. PR1345848
- After configuring and deleting the Ethernet loopback configuration, the interface goes down and does not come up. PR1353734
- On the QFX10000 line of switches, the analyzer does not mirror after adding the child member to an aggregated Ethernet interface. PR1417694
- The following error message is observed while performing NSSU: **syntax error: request-package-validate message.** PR1421378
- On the QFX5120 line of switches, one of the VCP ports of the throughput test result for most of the frame sizes is not close to 100 percent. PR1453709
- The **show interfaces xe-a/b/c** statement on a disabled or enabled configuration change displays fiber intermittently. PR1467509
- NSSU upgrade fails when there are multiple fpcs in the chassis NSSU upgrade group. PR1473624
- On the QFX5120-48T line of switches, convergence delay for the link-protected MPLS LSP is more than 50 minutes. PR1478584
- Observed 100 percent Layer 2 MAC scaling traffic loss in the QFX10002-60C line of switches after loading the EVPN-VXLAN collapsed profile configurations. PR1489753
- On the QFX5100 Virtual Chassis or Virtual Chassis fan, NSSU from the older Junos OS Release with Broadcom SDK 6.3.x to new Junos OS Release with Broadcom SDK 6.5.x might not work. PR1496765

Routing Protocols

- On the QFX5100 line of switches that does not run the QFX-5E codes (non TVP architecture), when image with Broadcom SDK upgrade (6.5.X) is installed, the CPU utilization might go up by around 5 percent. PR1534234

SEE ALSO

- What's New | 258
- What's Changed | 264
- Open Issues | 272
- Resolved Issues | 277
- Documentation Updates | 296
- Migration, Upgrade, and Downgrade Instructions | 297

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Release 20.1R3 for QFX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

Class of Service (CoS)

- DDoS violation on the QFX5200 line of switches is observed even after the received protocol packets are less than 10PPS. PR1381775

EVPN

- On the QFX10002 line of switches, the core link flaps and the BUM traffic loops. PR1492784
- The MAC address of the end-host gets wrongly programmed in the forwarding table after ESI failover. PR1584595

High Availability (HA) and Resiliency

- On the QFX5200-32C line of switches, the reboot time is degraded from 205 seconds in Junos OS Release 20.2R1 to 260 seconds in Junos OS Release 20.3. PR1511607

Infrastructure

- The following error message is observed during FTP: **ftpd[14105]: bl_init: connect failed for /var/run/blacklistd.sock (No such file or directory).** PR1315605

Interfaces and Chassis

- On the QFX5110 MC-LAG, flooding of the multicast packets for around 16 to 20 seconds is observed after disabling and enabling a member link of ICL after reboot. PR1422473
- ARP reply unicast packets might be flooded to all the interfaces in VLAN. PR1454764

Layer 2 Features

- On the QFX5000 Virtual Chassis, multicast traffic gets flooded even when the IGMP report times out. PR1431893
- On the QFX5000 line of switches, the following error message is reported in the **log: fpc0 Pools exhausted for Table:EGR_DVP_ATTRIBUTE_1**. PR1479826

- On the QFX5100 line of switches, fxpc CPU utilization is increased after the Broadcom SDK upgrade to 6.5.x from 5.3.x. PR1480132
- Traffic does not get load balanced by the QFX5000 line of devices over the ESI links with EVPN-VXLAN configured. PR1551543
- 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
- On the QFX5100 line of switches, traffic might be dropped in the Packet Forwarding Engine after change related to TPID when made in the dcd. PR1477156

Layer 2 Ethernet Services

- The DHCP DECLINE packets are not forwarded to the DHCP server when forward-only is set within dhcp-reply. PR1429456

Platform and Infrastructure

- On the QFX5100-48T-6Q line of switches, the port LEDs might not work. PR1317750
- On the QFX10000 line of switches, the source MAC and TTL values do not get updated for the routed multicast packets in EVPN-VXLAN. PR1346894
- The backup Routing Engine might crash after GRES occurs continuously for more than 10 times. PR1348806
- On the QFX10000 line of switches, the Aruba wireless access point (AP) heartbeat packets get dropped. As a result, the Aruba wireless AP cannot work. PR1352805
- USB upgrade of NOS image is not supported. PR1373900
- Due to the transient hardware condition, the single-bit error (SBE) events are corrected and have no operational impact. Those reported events had been disabled to prevent alarms and possibly unnecessary hardware replacements. PR1384435
- On Junos OS Release 18.4R1, intermittent traffic loss is observed with the RTG streams while flapping the RTG primary interface. PR1388082
- Unicast RPF in either the **Strict** mode or ICMP redirect does not work. PR1417546
- Memory leak is observed on the process l2ald when the rpd process is restarted. PR1435561
- On the QFX5200 line of switches, the ISSU might fail. PR1438690
- On the QFX5000 devices, the port qualifier is not supported. PR1440980
- On the QFX10000 line of switches, removal of the EVPN-VXLAN Layer 3 gateway on the IRB interface from the spine switches might cause traffic to be silently discarded. PR1446291

- On the QFX5000 line of switches, misleading ISSU logs are printed during the NSSU process even when the box does not perform ISSU. PR1451375
- Interface still sends mirrored traffic out even after it is removed from the RSPAN VLAN output.PR1452459
- Degradation of 9.51 percent with commit time and degradation of 12 percent with VLAN commit convergence are observed while comparing Junos OS Release 19.4DCB with Junos OS Release 19.3DCB. PR1457939
- On the QFX5110 line of switches, the VXLAN VNI (mcast) scaling causes traffic issue. PR1462548
- On the QFX5120-48T line of switches, finding discrepancy in the output of the **show chassis environment pem** command is observed in the backup member. PR1474520
- Interfaces are not detected on some of the ports when the 25-Gigabit Ethernet SFP is swapped and 10-Gigabit Ethernet SFP is inserted. PR1475574
- On the QFX5220 line of switches, the lo0 firewall filter might affect the Layer 3 forwarding traffic. PR1475620
- The **pfe_shm_vrf_hw_token_map_add** parameters are wrongly displayed as error message after loading base configuration. PR1480149
- On the QFX5100 Virtual Chassis or Virtual Chassis fan, NSSU from older Junos OS Release with Broadcom SDK 6.3.x to new Junos OS Release with Broadcom SDK 6.5.x might not work. PR1496765
- The QFX5110-48S-4c line of switches might have high 1 PPS output measurement error. PR1498739
- Kernel crash might occur after NSSU while performing GRES. PR1533874
- On the QFX5000 Virtual Chassis fan, traffic loss might be seen after swapping the primary and backup Routing Engines. PR1544353
- Need to move WRL7 to RCPL31 for the QFX-10-M and QFX-10-F line of switches. PR1547565
- Few LLDP sensor subscription do not work. PR1553534
- The MAC addresses learned in a Virtual Chassis might fail due to aging out in the MAC scaling environment. PR1558128
- While mapping analyzers to the channelized port, mirror might not work properly. PR1580473
- If the interface is newly added as the CE interface, the existing broadcast, unknown unicast, and multicast (BUM) traffic are looped. The loop prevention feature is designed to start working whenever a new CE interface is added by configuration. However, the existing BUM traffic arebe distributed to a new CE interface earlier before enabling the loop prevention feature. PR1493650
- Filter counter statistics verification fails when the received packets gets doubled. PR1590009
- On the QFX5100-48F-6Q switches, traffic loss is observed after de-activating and activating VLANs with VXLAN configurations. PR1592421
- On the QFX10000 line of swtiches, the active flows are not exported as expected. PR1442503

- The Layer 2 multicast traffic received on the VCP (Virtual Chassis port) ports might be dropped if igmp-snooping and STP/VSTP are enabled. PR1553159
- Upon the receipt of specific sequences of genuine packets destined to the device, the kernel crashes and restarts (vmcore). PR1557881
- The VCF might become unstable. PR1559172
- MAC addresses might not be relearned successfully after the MAC address age timeouts. PR1567723
- EVPN VXLAN CE interface with RSTP configured might cause LACP or BFD issues. PR1572504
- The WAN port links might not get brought down immediately during some abnormal type of line card reboot. PR1577315
- The Routing Engine kernel might crash due to logical child interface of the aggregated interface adding failure in the Junos kernel. PR1592456
- The existing ECMP route traffic might be dropped if you configure a static ECMP route with the same number of next hops as the existing ECMP route. PR1594573

Routing Protocols

- The dcpfe process generates core file after watchdog trigger caused by the failed MAC deletion notification. PR1371092
- On the QFX-5100 Virtual Chassis or Virtual Chassis Fan, the following error is observed in the hardware with the mini-PDT base configurations: **BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(), 128:l3 nh 6594 unintsall failed**. PR1407175
- The remaining BFD sessions of the aggregated Ethernet interface flap continuously if one of the BFD sessions is deleted. PR1516556
- The BFD sessions might flap continuously after disruptive switchover followed by GRES. PR1518106
- The rpd process might crash if next-hop self is used without using **extended-nexthop** and if the routing table has IPv4 routes with IPv6 nexthops. PR1582506
- The multi-hop BFD session might flap if you execute the **RSI (Request Support Information) collection** command. PR1589765

User Interface and Configuration

- The configuration under groups stanza is not inherited properly. PR1529989

Virtual Chassis

- On the QFX5000 Virtual Chassis, the DDoS violations that occur on the backup are not reported to the Routing Engine. PR1490552

SEE ALSO

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in Junos OS main and maintenance releases for QFX Series.

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

Resolved Issues: 20.1R3

Class of Service (CoS)

- Unable to configure policer with bandwidth-limit greater than 50G. PR1575049
- The buffer allocation for VCP ports might not get released in the Packet Forwarding Engine after physically moving the port location. PR1581187

EVPN

- Traffic might not get load balanced for multiple ESI/VTEP pairs with the underlay aggregated Ethernet interface between leaf and spine. PR1512253
- 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 multi-homed clients. PR1538117
- Policy with mac-filter-list might not work if you make changes that are unrelated to that policy and commit the changes in the EVPN scenario. PR1567623

Forwarding and Sampling

- The l2ald process might crash due to next-hop issue in the EVPN-MPLS. PR1548124
- Configuration archive transfer-on-commit fails on Junos OS Release 18.2R3-S6.5. PR1563641

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 observed after the MC-LAG interface. PR1562535
- Traffic loss might occur when you deactivate and activate member links of the ICL or ICCP interface. PR1542840
- New added MC-LAGs does not come up after the Routing Engine switchovers. PR1583547

Junos XML API and Scripting

- The **/var/run/scripts/** directory might be missing during bootup or upgrading the image. PR1543950

Layer 2 Features

- The dcpfe process might crash when the logical child interface continuously attaches and detaches. PR1543169
- Traffic might be forwarded incorrectly on an interface with VXLAN enabled and the **hold-time up xxx** statement configured. PR1550918
- On the QFX5120 line of switches, packets with VLAN ID 0 are dropped. PR1566850
- On the QFX5000 line of switches, software forwarded VXLAN decapsulated packets contains illegal length. PR1574435

- On the QFX5110-32Q line of switches, LACP does not come up in the Non-Oversubscribed mode for a set of ports. PR1563171
- In the OVSDB VXLAN scenario, inner VLAN tag 8 gets added unexpectedly into the encapsulated Ethernet header. PR1531319
- Traffic loop might occur in an MC-LAG scenario. PR1533301
- MAC programming issue occurs after deleting the IRB Layer 3 interface configuration from a VLAN configuration. PR1546179
- Packets received on a port that is in the **LACP Detached** and **Broadcom STP Blocked** states might get forwarded. PR1553570
- LACP gets into the **Detached** state when you delete VLAN on the aggregate interface configured on the SP style. PR1555862
- Traffic forwarding for VLAN 2 might not be correct when you remove a VLAN member from the ESI interface. PR1570446
- The dcpfe process crashes in the VXLAN scenario. PR1571170
- On the QFX5000 line of switches, DF might not forward the BUM traffic. PR1575976
- Traffic drop might occur 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 becomes offline for 120 seconds after the DHCP client sends the **DHCPINFORM** message in the DHCP relay scenario. PR1575740
- DHCP relay drops packets during the DHCP renewal process. PR1576417

Network Management and Monitoring

- Slow memory leakage might occur for the snmpd process. PR1575790

Platform and Infrastructure

- On the QFX10000 line of switches, 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
- The SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at the collector. PR1484322
- On the QFX5000 line of switches, multicast traffic loss is observed due to few multicast routes missing 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 line of switches, the cprod process timeout triggers high CPU utilization. PR1520956
- Packet drops might be seen with all commit events with 1G speed configured interface. PR1524614
- Traffic loss might be observed on the interfaces in a VXLAN environment. PR1524955
- On the QFX100002 line of switches, the firewall log incorrectly gets populated from the Packet Forwarding Engine. 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 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
- On the QFX5100-48T line of switches, interfaces are not created after 10G channel-speed is applied across the 48 to 53 ports. PR1538340
- ARP request might be dropped in a leaf device in a EVPN-VXLAN scenario. PR1539278
- The rpd memory leak might be observed on the backup Routing Engine due to the flapping of the link. PR1539601
- Unable to take RSI properly due to the authentication error. PR1539654
- FPC might not be recognized after power cycle (hard reboot). PR1540107
- The Packet Forwarding Engine might crash in the MPLS IPv6-tunneling scenario when the next hop changes. PR1540793
- The chip on FPC line card might crash when the system reboots. PR1545455
- OSPFv3 session might keep flapping and OSPFv3 hellos might be dropped in the host-path. PR1547032
- 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 40G 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
- The **action-shutdown** statement of storm control does not work for ARP broadcast packets. PR1552815
- Traffic might not pass 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 **then vlan** as the action. PR1556198
- Traffic storm might be caused by analyzer due to link flapping. PR1557274

- On the QFX5000 line of switches, the firewall filter might fail to work. PR1558320
- On the QFX5120 line of switches, amber LEDs are displayed for the fan modules after upgrading to Junos OS Release 20.2R1. PR1558407
- Few IPv6 ARP ND fails after loading the base configurations. PR1560161
- 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
- PTP lock status gets stuck at the **Acquiring** state instead of the **Phase Aligned** state. PR1561372
- On the QFX5000 line of switches, port mirroring might not work as expected. PR1562607
- On the QFX5120 line of switches, storm control with IRB interface might not work correctly. PR1564020
- On the QFX5100 line of switches, the following internal comment is displayed: **Placeholder for QFX platform configuration**. PR1567037
- On the QFX10002 line of switches, discrepancy in inet.1 vs Packet Forwarding Engine reported multicast routes. PR1567353
- On the QFX10000 line of switches, the firewall log is incorrectly populating from the Packet Forwarding Engine for IPv6 traffic. PR1569120
- On the QFX10008 chassis, the dcpfe process generates a core file. PR1572889
- On the QFX10000 line of switches, a high rate of 802.3X pause frames are sent out of the Interfaces. PR1575280
- The dcpfe process crashes while checking virtual tunnel-nh packet status. PR1580114
- On the QFX5120-32C line of switches, the following error is observed: **kern.ipc.maxpipekva exceeded; see tuning error**. PR1581192
- In the QFX10002-72Q line of switches, SNMP walk jnxOperatingEntry displays only two PSU even if four PSU are installed. PR1555852
- On the QFX5200 line of switches, the PRBS (Pseudo Random Binary Sequence) test fails for 100GbE interfaces with the default settings. PR1560086
- On the QFX10000 line of switches, the firewall filter logs are incorrectly populated the protocol 8847 entries. PR1582780
- When deleted aggregated Ethernet member(s) are not getting deleted (mirror trunk group) in the hardware for the analyzer input aggregated Ethernet. PR1589579
- The LCMD process might consume memory until all of the free memory available to VMHOST gets exhausted. PR1555386
- The dcpfe process might crash after committing the EVPN-VXLAN profile configuration and ARP resolution might fail causing traffic issues. PR1561588

- FPC might crash in a scaled-firewall configuration. PR1586817
- On the QFX10002 and QFX10008 line of switches, there might be traffic loss after FPC or system reboots. PR1487913
- The fxpc process might crash in an EVPN-VXLAN scenario. PR1504778
- On the QFX5110 with QSFP+40GE-IR4 line of switches, the unicast connectivity might break. PR1517601
- On the QFX5000 line of switches, the ECMP hash function might not take effect and the load balancing might not work. PR1523844
- On the QFX10000 line of switches, an enhancement to enable watchdog petting log on line cards is required. PR1527535
- The rpd process might crash due to memory leakage. PR1528550
- On the QFX5110-32Q line of switches, ports from 20 to 27 might flap when you insert the QSPF-40G transceiver into port 29 to 31. PR1535216
- On the QFX10000 line of switches, the Denial of Service (DoS) occurs upon receipt of DVMRP packets received on multi-homing ESI in VXLAN. PR1539194
- The **commit full** command might cause the guest VM to crash. PR1539434
- The aggregated Ethernet interface might flap after changing interface configurations. PR1542270
- Traffic loops if logical child interface gets added in the case of multihomed SP style in EVPN or VXLAN. PR1543966
- On the QFX10000 line of switches, the dcpfe process might crash. PR1546572
- On the QFX5000 line of switches, the static MAC on an interface might not work. PR1546655
- On the QFX10000 line of switches, ARP might not get resolved on the agrregated Ethernet interface. PR1546712
- LACP timeout issue might occur while polling for QSFP diagnostics. PR1549121
- The traffic are not load balanced properly in the EVPN overlay-ecmp setup. PR1550020
- The dcpfe process might crash due to chip SDK fault. PR1552645
- Traffic loss might occur on a VXLAN enabled VLAN. PR1554600
- The VGA might be down when you configure the IRB interface with multi VGA addresses. PR1555338
- Timestamp discrepancy might occur in the IPFIX packet flows exported. PR1558131
- The subscriber management infrastructure daemon (smid) process might get stuck at hundred percent. PR1559402
- On the QFX10000-60S-6Q line of switches, the line card takes more than 15 minutes to boot up after triggering the panic or watchdog reboot. PR1559725
- The VXLAN queue DDos violation and RARP packets flood might occur if receiving the RARP packets more than the supported DDoS bandwidth. PR1560243

- Sampled memory leak might occur when the analyzer is in the **Down** state. PR1564790
- Traffic loss might occur in the MC-LAG scenario. PR1565287
- The DF (Designated Forwarder) might not forward traffic. PR1567752
- On the QFX10002-60 line of switches, shutting down of one port causes another port to shutdown. PR1568294
- The BFD session flaps between the leaf and core during the spine reboot that causes other protocols to flap. PR1568615
- The dcpfe process might crash if the Type-5 tunnel fails to install for EVPN-VXLAN. PR1570136
- On the QFX10008 and QFX10016 line of switches, traffic loss might occur due to faulty FPC. PR1574779
- Port mirroring might not work when the analyzer output is a trunk interface. PR1575129
- On the QFX5000 line of switches, analyzer does not work. PR1576327
- The IS-IS packet might be corrupted on the provider edge device over the Layer 2 circuit tunnel. PR1580047
- The DHCP packets might get dropped if you apply the **dyn-dhcpv4_v6_trap** dynamic filter on the interface. PR1580352
- Multiple crashes with **toe_interrupt_errors** error message might occur. PR1593025

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 QFX 5100-48T-6Q Virtual Chassis or Virtual Chassis fan, the following error message is observed while copying the image to the Virtual Chassis fan member and trying to downgrade the image: **rcp for member 14, failed**. PR1486632
- Traffic might be silently discarded when the BGP route gets deleted, which is part of multipath. PR1514966
- The dcpfe process might crash while updating VRF for multicast routes during IRB uninit. PR1546745
- The BGP LU session might flap with AIGP-used scenario. PR1558102
- On the QFX5110-32Q lien of switches, 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
- The dcpfe process might crash when the size of the Local Bias Filter Bitmap string exceeds 256 characters. PR1568159
- The GRE egress traffic might not be forwarded between the different routing-instances. PR1573411

- The rpd crash might be observed after committing with static group 224.0.0.0 configured. PR1586631
- The dcpfe process might crash when any interface flaps. PR1579736
- Traffic might not be forwarding over the ECMP links in the EVPN VXLAN scenario. PR1533925
- The BFD sessions over IRB interface gets stuck in the **Init** state with FRR errors. PR1541851
- Multicast traffic with TTL 1 sent across VCP gets dropped. PR1543763
- BFD on the Layer 3 sub-interface of the ESI aggregated Ethernet interface might flap when an upstream underlay or overlay BGP flaps. PR1544982
- The rpd memory leak might occur in the BGP scenario. PR1547273
- On the QFX5000 line of switches, continuous traffic destined to a device configured with MC-LAG, that leads to nodes losing their control connection impacts traffic. PR1552877
- A filter could not be installed if the filter has a large scaled number of terms. PR1555337
- There might be traffic loss when the GRE interface flaps. PR1566428
- Memory leakage might occur in the MSDP scenario. PR1571906
- With IGMP snooping implemented, unexpected jitter issues might cause traffic loss. PR1583207

Resolved Issues: 20.1R2

Class of Service (CoS)

- PFC feature is not supported with QFX5120 Virtual Chassis due to chip limitation. PR1431895
- Traffic might be forwarded to the incorrect queue when a fixed classifier is used. PR1510365

EVPN

- The ESI of IRB interfaces does not update after autonomous-system number change if the interface is down. PR1482790
- The l2ald memory leakage might be observed in any EVPN scenario. PR1498023
- In the EVPN-VXLAN scenario, the l2ald process might crash in a rare condition. PR1501117
- The VXLAN function might be broken due to a timing issue. PR1502357
- Unable to create a new VTEP interface. PR1520078
- ARP table might not be updated in a race condition after performing VMotion or a network loop. PR1521526

Interfaces and Chassis

- The MC-LAG configuration-consistency ICL configuration might fail after committing some changes. PR1459201
- Traffic might get dropped as the next hop points to ICL even though the local MC-LAG is up. PR1486919

- MC-LAG consistency check fails if multiple IRB units are configured with same VRRP group. PR1488681
- Error message does not get generated while verifying the GRE limitation. PR1495543
- The dcpfe might crash when the ICL is disabled and then enabled. PR1525234

Layer 2 Ethernet Services

- Issues with the DHCPv6 relay processing confirm and reply packets are observed. PR1496220
- The MC-LAG might be down after disabling and then enabling the force-up configuration. PR1500758
- The aggregated Ethernet interface sometimes might not come up after the switch is rebooted. PR1505523

Layer 2 Features

- On the QFX5120 switches, the MAC learning might not work correctly. PR1441186
- On the QFX5120 switches, the third VLAN tag does not get pushed onto the stack. Instead, it gets swapped. PR1469149
- On the QFX5200 switches, the MAC learning rate is degraded by 88 percent. PR1494072
- Flow control is enabled in Packet Forwarding Engine irrespective of interface configuration and the fix causes a very small amount of packet loss when a parameter related to an interface such as "interface description" on any port is changed. PR1496766
- On the QFX5000 switches, traffic imbalance might be observed if hash-params is not configured. PR1514793
- The MAC address in the hardware table might become out of synchronization between the primary and member in Virtual Chassis after the MAC flaps. PR1521324

MPLS

- BGP session flaps between two directly connected BGP peers because of the wrong TCP-MSS in use. PR1493431

Platform and Infrastructure

- Traceroute monitor with MTR version v.69 shows a false 10 percent loss. PR1493824
- The following error message is generated while booting: **CMQFX: Error requesting SET BOOLEAN, illegal setting 66**. PR1385954
- The RIB installation or deletion time consumption is reduced. PR1421250
- SFP-LX10 stays down until autonegotiation is disabled. PR1423201
- The default logical interface on the channelized physical interface might not get created after ISSU or ISSR. PR1439358
- The PMTUD might not work for both IPv4 and IPv6 if the ingress Layer 3 interface is an IRB interface. PR1442587

- Members might stay disconnected from the QFX5120-32C/QFX5120-48T Virtual Chassis after a full-stack reboot. PR1453399
- Changing the VLAN name associated with the access ports might prevent the MAC addresses from being learned in the EVPN-VXLAN scenario. PR1454095
- On the QFX5000 line of switches, the dcpfe process crashes due to the usage of data that is not null getting terminated. PR1454527
- QFX5110 switch, the interface on QSFP-100GBASE-SR4 switch (made by Avago) cannot link up. PR1457266
- On the QFX5100 switches, the interface output counter is double-counted for self-generated traffic. PR1462748
- On the PTX10000 routers, FPCs might restart during runtime. PR1464119
- On the EX4600 device, traffic loss might be seen with framing errors or runts if MACsec is configured. PR1469663
- On the EX4600 device, DSCP marking might not work as expected if the fixed classifiers are applied to interfaces. PR1472771
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. PR1473610
- On the QFX5000 line of switches, the Layer 2 circuit might fail to communicate through VLAN 2. PR1474935
- The system might stop new MAC learning and impact the Layer 2 traffic forwarding. PR1475005
- sFlow does not work correctly if the received traffic goes out of more than one interface. PR1475082
- FPC major error is observed after the system boots up or the FPC restarts. PR1475851
- On the QFX10002-36Q/72Q switches, the following continuous error messages are logged on the device on getting adoption valid bit[8] asserted: **prds_ptc_wait_adoption_status: PECHIP[1] PTC[1]: timeout**. PR1477192
- Egress port mirroring might not work when the analyzer port and mirrored port belong to a different FPC. PR1477956
- SNMP Index in Packet Forwarding Engine reports as 0, causing SFLOW to report either IIF or OIF (not both) as 0 in sflow record data at collector. PR1484322
- VLAN creation failure might be observed with the scaled VLAN and Layer 3 configuration. PR1484964
- The dcpfe process might generate core files with the non-oversubscribed mode after SDK upgrade. PR1485854
- The 10GbE VCP ports will not be active in the QFX51XX and EX46XX Virtual Chassis scenario. PR1486002
- On the QFX5120 Virtual Chassis, the output of the show chassis alarm command displays incorrect PEM status after multiple GRES events. PR1486736

- QFX5100: If more than one UDF filter/term is configured, then only the first filter/term will be programmed in hardware. This is due to SDK 6.5.16 upgrade. PR1487679
- The queue statistics are not as expected after configuring the physical interface and logical interface shaping with the transmit rate and scheduler map. PR1488935
- After ISSU or ISSR, a port using SR4 or LR4 optics might not come up. PR1490799
- BFD sessions start to flap when the firewall filter in loopback0 is changed. PR1491575
- Junos OS: High CPU load due to receipt of specific multicast packets on Layer 2 interface (CVE-2020-1668). PR1491905
- Traffic loss could be observed in mixed Virtual Chassis setup of QFX5100 and EX4300. PR1493258
- Traffic loss might be seen in an MC-LAG scenario. PR1494507
- In the QFX5120 line of switches, the SNMP polling for the CPU utilization and state of the breakup-Routing Engine does not show in the two member Virtual Chassis. PR1495384
- Junos OS: PTX Series and QFX Series: Kernel routing table (KRT) queue stuck after J-Flow sampling of a malformed packet (CVE-2020-1679). PR1495788
- ARP might not get refreshed after timeout. PR1497209
- Virtual Chassis is not stable with 100GbE and 40GbE interfaces. PR1497563
- Outbound SSH connection flap or memory leak issue might be observed when pushing the configuration to the ephemeral database at him high rate. PR1497575
- On the QFX5210064C switches, the lcmd process generates a core file. PR1497947
- Traffic might get dropped if the aggregated Ethernet member interface is deleted or added, or an SFP transceiver of the aggregated Ethernet member interface is unplugged or plugged in. PR1497993
- The request-pfe-execute CLI command takes longer than 5 seconds to get a reply in Junos OS Release 18.4 for QFX5100. PR1498092
- On the QFX5210 switches, unexpected behavior for port LEDs lights is observed after the upgrade. PR1498175
- Inter-VNI and intra-VNI or VRF traffic is dropped between the CE devices when the interfaces connected between the TOR and multihomed PE devices are disabled. PR1498863
- On the QFX5100 and QFX5110 line of switches, the firewall filter might not get applied. PR1499647
- BFD sessions flap after deactivating or activating the aggregated Ethernet interface or executing GRES. PR1500798
- On the QFX5000 switches, ERPS might not work correctly. PR1500825
- The error message **mpls_extra NULL** might be seen during MPLS route add/change/delete operation. PR1502385
- The interface becomes physically down after changing to the FEC-none mode. PR1502959

- LLDP is not acquired when native-vlan-id and tagged VLAN-ID are the same on a port. PR1504354
- "Media type" in show interface command is displayed as "Fiber" for SFP-10G-T. PR1504630
- The DMA failure errors might be seen when the cache is full or flushes. PR1504856
- The l2cpd process might crash if the ERP configuration is added or removed, and the l2cpd process is restarted. PR1505710
- The archival function might fail in certain conditions. PR1507044
- The fxpc may crash and restart with an fxpc core file created while installing the image through ZTP. PR1508611
- Traffic might be affected on the QFX10002, QFX10008, and QFX10016 platform. PR1509220
- ARP replies might be flooded through the EVPN-VXLAN network as unknown unicast ARP reply. PR1510329
- The output VLAN push might not work. PR1510629
- Multicast traffic loss is observed because of few missing multicast routes in the spine node. PR1510794
- The QFX10000-36Q line card used on QFX10008 and QFX10016 platforms may fail to detect any QSFP. PR1511155
- In the VXLAN configuration, the firewall filters might not be loaded into the TCAM with the following message due to TCAM overflow after upgrading to Releases 18.1R3-S1, 18.2R1, and later : DFWE ERROR DFW: Cannot program filter. PR1514710
- The routes update might fail upon the HMC memory issue and traffic impact might be seen. PR1515092
- The100GbE AOC non-breakout port might be auto-channelized to another speed. PR1515487
- The MAC learning might not work properly after multiple MTU changes on the access port in the VXLAN scenario. PR1516653
- The dcpfe (PFE) process might crash due to memory leak. PR1517030
- The vgd process might generate a core file when the OVSDB server restarts. PR1518807
- Traffic forwarding might be affected when adding, removing, or modifying the VLAN or VNI configurations such as vlan-id and vni-id, and the ingress-replication configuration. PR1519019
- Output interface index in an sFlow packet is zero when transit traffic is observed on the IRB interface with VRRP enabled. PR1521732
- On the QFX10002, QFX10008, and QFX10016 switches, the following error message is observed during specific steps while clearing and loading the scaled configuration again: PRDS_SLU_SAL:jprds_slu_sal_update_lrncnt(),1379: jprds_slu_sal_update_lrncnt call failed. PR1522852
- Sampling, with the rate limiter command enabled, crosses the sample rate 65,535. PR1525589

- The MPLS EXP classifier might not work on QFX10000 platforms. PR1531095
- High rate of ARP or NS packets might be observed between a device that runs Junos OS and host when the device that runs Junos OS receives an ARP or NS packet on an interface in transition. PR1534796

Routing Policy and Firewall Filters

- The policy configuration might be mismatched between rpd and mgd when **deactivate policy-options prefix-list** is involved in configuration sequence. PR1523891

Routing Protocols

- Flows do not fall back to a single link when the inactivity-interval is set higher than the IFG. PR1471729
- The MUX state in the LACP interface does not go to the Collecting and Distributing state and remains in the Attached state after enabling the aggregated Ethernet interface. PR1484523
- The FPC process goes to the NotPresent state after upgrading the QFX5100 Virtual Chassis or Virtual Chassis Fan. PR1485612
- On QFX 5100-48T-6Q with Virtual Chassis or Virtual Chassis fan, system upgrade/ installation might fail. PR1486632
- CPU port queue gets full due to excessive pause frames being received on interfaces; this causes control packets from the CPU to all ports to be dropped. PR1487707
- The BGP route target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743
- The rpd process generates core files at **rt_nh_resolve_add_gen** in **../../../../../../../src/junos/usr.sbin/rpd/lib/rt/rt_resolve_ind.c:** with the evpn-dhcp configurations. PR1494005
- EX4300-MP/EX4600/QFX5000 Series: High CPU load due to receipt of specific layer 2 frames in EVPN-VXLAN deployment. (CVE-2020-1687) & High CPU load due to receipt of specific Layer 2 frames when deployed in a Virtual Chassis configuration (CVE-2020-1689). PR1495890
- Firewall filter does not work in certain conditions in a Virtual Chassis setup. PR1497133
- Traffic drop might be observed after modifying the FBF firewall filter. PR1499918
- Scale of filters with egress-to-ingress command is enabled. PR1514570
- The rpd might report 100% CPU usage with BGP route damping enabled. PR1514635
- Firewall "sample" configuration gives the warning as unsupported on QFX10002-36Q and does not work. PR1521763
- On the QFX5000 line of switches, the fxpc process might crash if the VXLAN interface flaps. PR1528490

User Interface and Configuration

- The version information under the configuration changes from Junos OS Release 19.1 onward. PR1457602

Resolved Issues: 20.1R1

Class of Service (CoS)

- Shaping does not work after the reboot if **shaping-rate** is configured. PR1432078
- The traffic is placed in the network-control queue on an extended port even if it comes in with a different DSCP marking. PR1433252
- On QFX5120 switches, when you move unicast traffic to a multicast queue through an MF classifer, the **show interface queue** command does not display any status. PR1459281

EVPN

- The rpd might crash with EVPN-related configuration changes in a static VXLAN to MPLS stitching scenario. PR1467309

Forwarding and Sampling

- Type 1 ESI/AD route might not be generated locally on an EVPN PE device in the **all-active** mode. PR1464778

General Routing

- On QFX5100 Virtual Chassis, **MacDrainTimeOut and bcm_port_update failed: Internal error** is observed. PR1284590
- The **show chassis errors active detail** command is not supported on QFX5000 platforms. PR1386255
- The 10-Gigabit Ethernet fiber interfaces might flap frequently when they are connected to other vendor's switch. PR1409448
- The optic comes with Tx enabled by default. As the port is administratively disabled, the port is stopped but as the port has not been started, it does not disable Tx. PR1411015
- Part of routes could not be provided into the Packet Forwarding Engine when both IPv4 and IPv6 are used. PR1412873
- The **show interface** command shows **Media type: Fiber** on QFX5100-48T switches running "QFX 5e Series" image. PR1419732
- Ports might get incorrectly channelized if they are channelized to 10-Gbps and they are again channelized to 10-Gbps. PR1423496
- CoS rewrite rules applied under an aggregated Ethernet interface might not take effect after nonstop software upgrade (NSSU). PR1430173
- The l2cpd process might crash and generate a core file when interfaces flap. PR1431355

- The FPC might crash when a firewall filter is modified. PR1432116
- When you plug in an unsupported SFP-T module, the line card might crash. PR1432809
- BGP neighborship might not come up if the MACsec feature is configured. PR1438143
- QFX5100 Virtual Chassis does not come up after you replace a Virtual Chassis port fiber connection with a DAC cable. PR1440062
- MAC addresses learned on RTG might not be aged out after a Virtual Chassis member is rebooted. PR1440574
- Packet loss might be seen if IPoIP or MPLS-over-UDP dynamic tunnels are configured with ECMP. PR1446132
- On QFX5100 Virtual Chassis, a cyclic redundancy check (CRC) error might be seen on the Virtual Chassis Port (VCP). PR1449406
- Except one aggregated Ethernet member link, the other links do not send out sFlow sample packets for ingress traffic. PR1449568
- The em0 route might be rejected after the em0 interface is disabled and then enabled. PR1449897
- FPC does not restart immediately after rebooting the system. This might cause packet loss. PR1449977
- On QFX10000 platforms, CoS classification does not work. PR1450265
- The l2ald and eventd process are hogging 100 percent after the **clear ethernet-switching table** command is issued. PR1452738
- The classifier configuration does not get applied to the interface in an EVPN-VXLAN environment. PR1453512
- The **show chassis led** command shows incorrect status. PR1453821
- On QFX5100 Virtual chassis, VGD process hogs the CPU without the **switch-options vtep-source-interface lo0.0** configuration. PR1454014
- On QFX5110 Virtual Chassis, master FPC might come up in master state again after reboot instead of backup. PR1454343
- On QFX5000 platform, the dcpfe process crashes because usage of data which is not NULL is terminated. PR1454527
- On QFX10002-60C EVPN-VXLAN, the MAC+IP count is shown as zero. PR1454603
- On QFX5120 switches, untagged hosts ARP/NS requests connected on **encapsulation ethernet-bridge** interface are not being resolved. PR1454804
- You might not be able to apply a firewall filter to a particular Virtual Chassis or Virtual Chassis Fabric member as TCAM is running out of space. PR1455177
- In a 16+ member QFX5100 Virtual Chassis Fabric, the **FROM** column under the **show system users** command output reports feb0, feb1, feb2, and feb3 for fpc16, fpc17, fpc18, and fpc19, respectively. PR1455201

- The priority-based flow control (PFC) feature does not work on the QFX10000 line of switches. PR1455309
- The cosd crash might be observed if the **forwarding-class-set** is directly applied on the child interface of an aggregated Ethernet interface. PR1455357
- Link-up delay and traffic drop might be seen on mixed service provider Layer 2/Layer 3 and enterprise style Layer 2 type configurations. PR1456336
- The Packet Forwarding Engine process might crash after Routing Engine switchover on QFX10000 platforms. PR1457414
- Overtemperature SNMP trap messages are displayed after an update even though the temperatures are within the system thresholds. PR1457456
- On QFX5110 switches, port 51 has one LED blinking amber continuously. PR1457516
- On QFX5210 switches, the LED does not light on port 64 and 65 after the switch is upgraded to Junos OS Release 19.2R1. PR1458514
- The command **show dynamic-tunnels database** does not show **v6 mapped** next-hop flag for 6PE routes that have labels. PR1458634
- The BPDU packet might be looped between leaf DF switch and non-DF switch and causes traffic blocking. PR1458929
- On QFX5200 switches, DHCPv6 LDRA relay bounded count is not as expected after DHCP is configured. PR1459499
- The fxpc process might crash because the BGP IPv6 session flaps. PR1459759
- The **forwarding** option is missed in routing instance type. PR1460181
- The **accept-source-mac** feature with VXLAN is not working on QFX5000 platforms. PR1460885
- The statement **show forwarding-options enhanced-hash-key** is not supported on QFX10000 platforms. PR1462519
- The entPhysicalTable MIB is not fetching expected data on QFX10002-72Q or QFX10002-36Q platforms. PR1462582
- The fxpc process might generate core files when changing MTU in a VXLAN scenario with firewall filters applied on QFX5000 platforms. PR1462594
- On QFX5100 Virtual Chassis or Virtual Chassis Fabric, you observe the **BRCM-VIRTUAL,brcm_vxlan_walk_svp(),6916:Failed to find L2-iff for ifl:** error while cleaning up EVPN-VXLAN configurations with mini-PDT base configurations. PR1463939
- On PTX10000, the FPC might restart during runtime. PR1464119
- On QFX10000 platforms, the interface might not come up on FPC restart. PR1464650
- QFX5100-24Q: Unable to apply DSCP rewrite to firewall filter to a Layer 3 subinterface (for example, xe-0/0/0.100). PR1464883

- PEM is not present spontaneously on QFX5210. PR1465183
- On QFX5100-48T switches, a 10-Gigabit Ethernet interface might not come up or negotiate at speed 1-Gbps when connected with BRCM 10G/GbE 2+2P 57800-t rNDC. PR1465196
- The QSFP-100G-PSM4 could not be correctly identified on QFX5200 or QFX5110 platforms. PR1465214
- The physical interface of an aggregated Ethernet might take time to come up after disabling or enabling it. PR1465302
- Junos OS exhibits inconsistent fan and power supply numbering on White Boxs (-O and -OZ) in Release 19.2R1. PR1465327
- In a Virtual Chassis scenario, the broadcast and multicast traffic might be dropped over an IRB or a LAG interface. PR1466423
- BGP open messages with specific types of BGP optional capabilities causing BMP messages not to be encoded correctly when sent to the BMP collector. PR1466477
- On QFX10000 platforms, EBUF parity interrupt is not seen. PR1466532
- IPv6 traffic over Layer 3 VPN might fail. PR1466659
- Slow packet drops might be seen on QFX5000 platforms. PR1466770
- EPR iCRC errors in QFX10000 platforms might cause protocols to be down. PR1466810
- A few of the DHCPvX INFORM messages, specific to a particular VLAN, are not receiving any ACK from server. PR1467182
- Ingress drops to be included at the CLI from interface statistics and added to InDiscards. PR1468033
- Optics measurements might not be streamed for interfaces of a PIC over JTI. PR1468435
- MAC address might not be learned on a new extended port after VMotion in a Junos fusion for data center environment. PR1468732
- QFX5000 platform is looping the IP routed packet through IS-IS or MPLS. PR1469998
- Incorrect counter values are observed for the arrival rate and peak rate for DDoS commands. PR1470385
- On QFX5100 and EX4300 mixed-mode Virtual Chassis, unable to configure 10-Mbps speed on the Gigabit Ethernet interface. PR1471216
- In a VXLAN scenario on QFX10000 platforms, when a VTEP source interface is configured in multiple routing instances, traffic loss might occur. PR1471465
- On QFX5000 platforms, egress PACL size is half. PR1472206
- The shaping of CoS does not work after reboot. PR1472223
- The detached interface in a LAG might process the xSTP BPDUs. PR1473313
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. PR1473685
- On QFX5000 platforms in an EVPN-VXLAN scenario, continuous log messages might be observed. PR1474545

- Layer 2 circuit might fail to communicate via VLAN 2 on QFX5000 platforms. PR1474935
- DAC cables are not being properly detected in the Packet Forwarding Engine on QFX5200 switches. PR1475249
- QFX5000 leaf device might fail to forward the traffic in a multicast environment with VXLAN. PR1475430
- QFX Series platform generates the **invalid PFE PG counter pairs to copy, src 0xfffff80, dst 0** message. PR1476829
- On QFX10002-36Q and QFX10002-72Q switches, generating continuous **prds_ptc_wait_adoption_status: PECHIP[1] PTC[1]: timeout on getting adoption valid bit[8] asserted** error logs on the device. PR1477192
- The remaining interface might be still in downstate even the number of channelized interfaces is no more than five. PR1480480
- ARP request packets for unknown hosts might get dropped in a remote PE in an EVPN-VXLAN scenario. PR1480776
- On QFX10000 and QFX5000 Series switches with SP style configuration, BUM traffic incorrectly get blocked, while you disable or enable different logical interfaces. PR1482202
- After an ISSU or an ISSR, a port using SR4 or LR4 optics might not come up. PR1490799

High Availability (HA) and Resiliency

- Unified ISSU is not supported on QFX5000 platforms. PR1472183

Interfaces and Chassis

- VRRPv6 state is flapping with init and idle states after configuring **vlan-tagging**. PR1445370
- Traffic might be forwarded to incorrect interfaces in an MC-LAG scenario. PR1465077
- On a QFX Series platform, VRRPv3 MIBs are not working to poll VRRPv6-related objects. PR1467649
- Executing commit might become unresponsive due to a stuck dcd process. PR1470622
- Commit error is not thrown when a member link is added to multiple aggregation groups with different interface-specific options. PR1475634

Junos Fusion Enterprise

- Loop detection might not work on extended ports in Junos fusion for enterprise scenarios. PR1460209

Junos Fusion Satellite Software

- In Junos fusion for enterprise, dpd might crash on satellite devices running SNOS. PR1460607

Layer 2 Ethernet Services

- In an EVPN-VXLAN ERB scenario, **dhcp relay-source lo0.1** is not used when enabled with anycast legacy IRB. PR1455076

- Member links state might be asychronized on a connection between PE and CE devices in an EVPN A/A scenario. PR1463791

Layer 2 Features

- On QFX5100 switches, storm control configuration might be disabled for the interface. PR1354889
- Physical layer and MAC/ARP learning might not work for copper base SFP-T transceivers on QFX5100 and QFX5110. PR1437577
- The LLDP function might fail when a Juniper device connects to a non-Juniper device. PR1462171
- A few MAC addresses might be missing from the software MAC table on QFX5000 platforms. PR1467466
- After rebooting, an FXPC core file might be seen when committing the configuration. PR1467763
- Ingress traffic might be silently dropped if the underlying interface flaps in an EVPN-VXLAN scenario. PR1469596
- Traffic might be affected if composite next hop is enabled. PR1474142

MPLS

- On QFX10002 switches, the **show mpls static-lsp | display xml** command produces invalid XML. PR1469378
- Traffic might silently dropped and discarded on PE when CE sends traffic to PE and the destination is resolved with two LSPs through one upstream interface. PR1475395
- MPLS LDP ping or traceroute fails over QFX5100 as transit PHP node. PR1477301

Platform and Infrastructure

- The stylesheet language alternative syntax (SLAX) script might be lost after upgrading software. PR1479803

Routing Protocols

- In a scaled setup, when the host table is full and the host entries are installed in the LPM table, OSPF sessions might take more time to come up. PR1358289
- Invalid VRRP mastership election on QFX5110 Virtual Chassis peers. PR1367439
- Host-destined packets with **filter log** action might not reach the Routing Engine if log/syslog is enabled. PR1379718
- On QFX5100, BGP IPv4 or IPv6 convergence and RIB install or delete time degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. PR1414121
- PIM (S, G) joins can cause MSDP to incorrectly announce source active messages in some cases. PR1443713
- CRC errors might be seen on QFX5100 Virtual Chassis. PR1444845
- The core file might be generated when you add or remove EVPN Type-5 routing instance. PR1455547

- On QFX5000 platforms, egress port for ARP entry in the Packet Forwarding Engine is not modified from the VTEP to the local ESI port, after the device boots up.PR1460688
- On QFX5100 Virtual Chassis or Virtual Chassis Fabric, the **brcm_ipmc_route_counter_delete:3900Multicast stat destroy failed (-10:Operation still running)** error is observed after unified ISSU with Mini-PDT base configurations. PR1460791
- The **other querier present interval** timer cannot be changed in an IGMP/MLD snooping scenario. PR1461590
- When IRB is deleted on the Layer 3 gateway, the IRB interface does not get removed from the Packet Forwarding Engine and it results in traffic drop in IRB MAC address. PR1463092
- The mcsnoopd crash might be seen if one BD/VLAN is configured as part of EVPN and it has any multicast router interfaces (static/dynamic). PR1468737
- Traffic might not be forwarded over an ECMP link in an EVPN-VXLAN scenario. PR1475819
- ARP packets are always sent to CPU regardless of whether the **storm-control** is activated. PR1476708
- GRE transit traffic is not forwarded in a VRRP scenario. PR1477073

SEE ALSO

What's New | **258**

What's Changed | **264**

Known Limitations | **270**

Open Issues | **272**

Documentation Updates | **296**

Migration, Upgrade, and Downgrade Instructions | **297**

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 20.1R3 documentation for the QFX Series.

Dynamic Host Configuration Protocol (DHCP)

- **Introducing DHCP User Guide**—Starting in Junos OS Release 20.1R1, we are introducing the DHCP User Guide for Junos OS routing, switching, and security platforms. This guide provides basic configuration details for your Junos OS device as DHCP Server, DHCP client, and DHCP relay agent.

  [See DHCP User Guide.]

SEE ALSO

- What's New | 258
- What's Changed | 264
- Known Limitations | 270
- Open Issues | 272
- Resolved Issues | 277
- Migration, Upgrade, and Downgrade Instructions | 297

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 **20.1** 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 20.1 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-x86-64-20.1-R2.n-secure-signed.tgz reboot**

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

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

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

Rebooting occurs only if the upgrade is successful.

> **NOTE:** After you install a Junos OS Release 20.1 **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-x86-64-20.1R3.9.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-x86-64-20.1R3.9.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-x86-64-20.1R3.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-x86-64-20.1R3.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-x86-64-20.1R3.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-x86-64-20.1R3.n-secure-signed.tgz
    ```

    For more information about the **request system software add** command, see the CLI Explorer.

15. Reboot the Routing Engine using the **request system reboot** command:

```
user@switch> request system reboot
```

> **NOTE:** You must reboot to load the new installation of Junos OS on the switch.
>
> To abort the installation, do not reboot your system. Instead, finish the installation and then issue the **request system software delete jinstall *<package-name>*** command. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not send traffic.

16. Log in and issue the **show version** command to verify the version of the software installed.

17. Transfer routing control back to the master Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the CLI Explorer.

18. Verify that the master Routing Engine (slot 0) is indeed the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
  Slot 0:
    Current state                  Master
    Election priority              Master (default)
outing Engine status:
  Slot 1:
    Current state                  Backup
    Election priority              Backup (default)
```

Performing a Unified ISSU

You can use unified ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.

NOTE: Unified ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- Preparing the Switch for Software Installation on page 308
- Upgrading the Software Using Unified ISSU on page 308

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-x86-64-20.1R3.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

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

provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

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

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

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

SEE ALSO

- What's New | 258
- What's Changed | 264
- Known Limitations | 270
- Open Issues | 272
- Resolved Issues | 277
- Documentation Updates | 296

Junos OS Release Notes for SRX Series

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 for the SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

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

Release 20.1R3 New and Changed Features

There are no new features in Junos OS Release 20.1R3 for the SRX Series devices.

Release 20.1R2 New and Changed Features

There are no new features in Junos OS Release 20.1R2 for the SRX Series devices.

Release 20.1R1 New and Changed Features

Application Security

- **Custom application enhancements (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.1R1, we've enhanced the custom applications signature functionality by providing a new set of applications and contexts.

  Application identification allows you to create custom application signatures to detect applications specific to your network environment. You can create custom application signatures for applications based on ICMP, IP protocol, IP address, and Layer 7 or TCP/UDP stream. While configuring the custom application signatures, you must specify the context values that the device can use to match the patterns in the application traffic.

Custom application signature contexts are part of application signature package. You must download and install the latest application signature package version 3248 or later to use new contexts for custom application signatures.

[See Custom Application Signatures for Application Identification.]

- **Default mechanism to forward the traffic through APBR rule (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS 20.1R1, you can configure a APBR rule by specifying the dynamic application match criteria with any keyword. This provides a default mechanism to forward the traffic to a specific next-hop device or to a destination if the traffic matches any dynamic application.

  [See Advanced Policy-Based Routing.]

- **AppQoE support for granular APBR rules (NFX Seris, SRX Series, and vSRX)**—Starting in Junos OS Release 20.1R1, AppQoE utilizes the granular rule matching functionality of advanced policy-based routing (APBR) for better quality of experience (QoE) for the application traffic.

  In Junos OS Release 18.2R1, APBR supported configuring policies by defining source addresses, destination addresses, and applications as match conditions. After a successful match, the configured APBR profile is applied as an application services for the session. In this release, AppQoE leverages the APBR enhancement and selects the best possible link for the application traffic as sent by APBR to meet the performance requirements specified in SLA.

  [See Application Quality of Experience.]

Authentication and Access Control

- **Support for UPN as user identity (SRX Series)**—Starting in Junos OS Release 20.1R1, you can use User Principal Name (UPN) as logon name in firewall-authentication, which is working as a captive portal for JIMS or user-firewall.

  You can use UPN as logon name along with *cn* or *sAMAccountName* at the same time. UPN can be used instead of *sAMAccountName* to authenticate a user.

  Even if user uses UPN as logon name, firewall authentication pushes *sAMAccountName* (mapping to the UPN) to user ID rather than pushing the UPN.

  Firewall-authentication pushes both UPN and *sAMAccountName* (mapping to the UPN) to JIMS.

  [See Understanding Advanced Query Feature for Obtaining User Identity Information from JIMS.]

- **Trusted Platform Module (TPM) to bind secrets (SRX5400, SRX5600, and SRX5800)**—Starting with Junos OS Release 20.1R1, we've introduced the TPM support on the SRX5000 line of devices with SRX5K-RE3-128G Routing Engine (RE3). The TPM chip is enabled by default to make use of TPM functionality.

  When TPM is activated, it protects the private keys stored in Junos OS.

  [See Using Trusted Platform Module to Bind Secrets on SRX Series Devices.]

Flow-Based and Packet-Based Processing

- **Support of IPFIX formatting and Chassis Cluster for SRX J-Flow functionality (SRX300, SRX320, SRX340, SRX345, and SRX550HM)** —Starting with Junos OS Release 20.1R1, you can configure Chassis Cluster and define an IPFIX flow record template suitable for IPv4 traffic or IPv6 traffic. IPFIX is an enhanced version of J-flow version 9 template. Using IPFIX, you can collect a set of sampled flows and send the record to a specified host.

  See [Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, and SRX devices.]

- **Support service inspection for pass-through IP-IP and GRE tunnel in TAP mode (SRX300, SRX320, SRX340, SRX345, SRX1500, SRX4100, and SRX4200)**—Starting in Junos OS Release 20.1R1, TAP mode inspects IP-IP and GRE inner tunnel traffic by de-encapsulating the outer and inner IP header (up to two levels) to create flow sessions. You can configure up to eight TAP interfaces on an SRX Series device.

  [See TAP Mode for Flow Sessions, and forwarding-options.]

GPRS

- **Increase in GTP scale for IoT and roaming firewall applications (SRX1500, SRX4100, SRX4200, and vSRX)**—Starting in Junos OS Release 20.1R1, in addition to the existing support on SRX5400, SRX5600, SRX5800, and SRX4600, to enable the Internet of Things (IoT) and roaming firewall use cases, the GTP tunnel scale is increased for the following SRX Series devices:

  - SRX1500: 204,800 to 1,024,000
  - SRX4100: 409,600 to 4,096,000
  - SRX4200: 819,200 to 4,096,000

  For vSRX instances, the number of tunnels supported depends on the available system memory.

[See Understanding Policy-Based GTP.]

Hardware

- **SRX380 Services Gateway**—The SRX380 Services Gateway is a high performance and all-in-one networking device, which consolidates routing, switching, and security. With next-generation firewall features and advanced threat mitigation capabilities, the SRX380 device provides cost-effective and secure connectivity across distributed enterprise locations. A 1U form factor model with a 16-core MIPS processor and 4-GB DDR4 RAM, the SRX380 device supports up to 10-Gbps firewall performance.

  The SRX380 device has an integrated 100-GB SSD and provides high port density with 16 on-board PoE-enabled 1-Gigabit Ethernet RJ-45 ports and 4 10-Gigabit Ethernet SFP+ ports. All the ports support AES-256 MACsec encryption. The SRX380 device has dual AC power supplies and supports up to four Mini-PIMs.

  The SRX380 supports the same features as those supported on the existing SRX300 line of services gateways. For the complete list of features supported on the SRX380, see Feature Explorer.

  [See SRX380 Services Gateway Overview.]

Interfaces and Chassis

- **Support for new show | display set CLI commands (ACX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the following new **show** commands have been introduced:

  - **show | display set explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the top level of the hierarchy.
  - **show | display set relative explicit**—Display explicitly, as a series of commands, all the configurations that the system internally creates when you configure certain statements from the current hierarchy level.

  [See show | display set and show | display set relative.]

Intrusion Detection and Prevention

- **HTTP X-Forwarded-For header support in IDP (SRX Series)**—Starting from Junos OS Release 20.1R1, we've introduced the **log-xff-header** option to record the x-forward-for header (xff-header) information. When this option is enabled. During the traffic flow, IDP saves the source IP addresses (IPv4 or IPv6) from the contexts for HTTP and SMTP traffics and displays in attack logs.

  The xff-header is not processed unless its enabled through sensor-configuration.

  - To enable the xff-header, use the **set security idp sensor-configuration global log-xff-header** command.
  - To disable the xff-header, use the **delete security idp sensor-configuration global log-xff-header** command.

Previously, when you access internet, to lessen the external bandwidth the servers used transparent proxies. It was difficult to identify the originating source IP address as the proxy server converted it into an anonymous source IP address.

[See Understanding Multiple IDP Detector Support.]

Juniper Sky ATP

- **Juniper Sky ATP support for disabling standard Juniper C&C and URL feeds**—Starting in Junos OS Release 20.1R1, you can disable standard Juniper command and control (C&C) and URL feeds on SRX Series devices. Disabling the Juniper C&C and URL feeds helps to free the resources on SRX Series devices and makes the resources available for loading custom feeds. Use the **set services security-intelligence disable-global-feed (all | feed name *feed-name*)** command to disable the feeds. To enable the feeds, use the **delete services security-intelligence disable-global-feed (all | feed name *feed-name*)** command.

  [See set services security-intelligence and show services security-intelligence category summary.]

Junos OS XML API and Scripting

- **The jcs:load-configuration template supports loading the rescue configuration (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **jcs:load-configuration** template supports the **rescue** parameter to load and commit the rescue configuration on a device. SLAX and XSLT scripts can call the **jcs:load-configuration** template with the **rescue** parameter set to **"rescue"** to replace the active configuration with the rescue configuration.

  [See Changing the Configuration Using SLAX and XSLT Scripts and jcs:load-configuration Template.]

J-Web

- **J-Web supports SRX380 device**—Starting in Junos OS Release 20.1R1, you can use J-Web to manage your SRX380 device. Additionally, you can also:
  - Monitor wireless LAN setting of the supported Wi-Fi Mini-PIM: Monitor > Wireless LAN.
  - View power statistics information using the new Power Budget Statistics tab: Monitor > Chassis Information > Chassis Component Details.

    NOTE: You can view the power statistics information only when the device is in standalone mode.

  - Configure wireless LAN setting of the supported Wi-Fi Mini-PIM: Configure > Wireless LAN > Settings.
  - Configure redundant power supply for power management using the new Redundant PSU menu: Configure > Basic Settings.

[See Dashboard Overview, Monitor Wireless LAN, and About the Settings Page.]

Network Management and Monitoring

- **SNMP support to export statistics of user firewall (SRX Series and vSRX)**—Starting in Junos OS Release 20.1R1, the new MIB jnxUserFirewalls OID is introduced to expose statistics of user firewall identity-management counters to network monitoring tools supporting SNMP.

  [See Enterprise-Specific SNMP MIBs Supported by Junos OS.]

- **SNMP support to monitor Express Path status (SRX4600, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.1R1, the new SNMP MIB jnxJsFlowSofSummary is introduced to improve the Express Path mode (formerly known as services offloading) session status using CLI monitoring and traffic logging. The jnxJsFlowSofSummary MIB Provides the total number of Express Path sessions in use and total number of packets processed so far in the logical system.

  [See Enterprise-Specific SNMP MIBs Supported by Junos OS.]

- **Enhanced PKI traps, log notifications, and SNMP for IPsec VPN (MX Series with USF and the SRX5000 line of devices with SPC3 card)**—Starting in Junos OS Release 20.1R1, you can enable the peer down and IPsec tunnel down traps and configure the certificate authority (CA) and local certificate traps. We've enhanced the existing IPsec VPN flow monitor MIB **jnxIpSecFlowMonMIB** to support the global data plane, active IKE SA, active IPsec SA, and active peer statistics for tunnels using IKEv2. We've also enhanced the output of the **show security ike stats** command to add additional options **(<brief> | <detail>)**. Use the **clear security ike stats** command to clear the IKEv2 statistic counters.

  [See Configure the Certificate Expiration Trap, Enterprise-Specific SNMP MIBs Supported by Junos OS, Enable Peer Down and IPsec Tunnel Down Traps, trap (Security PKI), trap (Security IKE), clear security ike stats, show security ike stats, show security ipsec statistics, show security ike security-associations, and show security ike active-peer.]

Port Security

- **Media Access Control Security (MACsec) support (SRX380)**—SRX380 supports MACsec in on all 16 1GbE ports and all four 10GbE ports. MACsec is an industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. The supported cipher suites are GCM-AES-256 and GCM-AES-128. Only static CAK mode is supported.

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

Security

- **Support for security policy reports (SRX Series and vSRX)**—Starting in Junos OS Release 20.1R1, you can use the **show security policy-report** command to display detailed security policy reports.

  Optimizing security policies ensure that the policies are efficient. Over time, policies become disorganized and hence ineffective. You can use the **show security policy-report** command to notify end users when you create new policies or change existing policies that adversely affect other security policies.

You can use the **report-skip** command at the **[edit security policies from-zone *zone-name* to-zone *zone-name* policy *policy-name*]** hierarchy level to exclude the policy from the policy analysis and prevent it from appearing in any future report.

[See show security policy-report and report-skip.]

- **Support to clear DNS cache if DNS error responses are received (SRX Series and vSRX)**—Starting in Junos OS Release 20.1R1, you can clear the DNS cache entry IP list when DNS error responses are received. We have introduced a new command, **dns-cache** under the **[edit security policies]** hierarchy level, to configure the security policy DNS cache behavior.

  [See dns-cache.]

System Management

- **Restrict option under NTP configuration is now visible (ACX Series, QFX Series, MX Series, PTX Series, and SRX Series)**—Starting in Junos OS Release 20.1R1, the **noquery** command under the **restrict** hierarchy is now available and can be configured with a mask address. The **noquery** command is used to restrict ntpq and ntpdc queries coming from hosts and subnets.

  [See Configuring NTP Access Restrictions for a Specific Address.]

Tenant Systems and Logical Systems

- **ICAP service redirect support for tenant systems (SRX Series and vSRX)**—You can prevent data loss from your network by employing Internet Content Adaptation Protocol (ICAP) redirect services. Starting in Junos OS Release 20.1R1, you can enable ICAP at the tenant system level, and you can view/clear the ICAP services redirect status and statistics at the tenant systems level.

  In addition, we've introduced the **X-Client-IP**, **X-Server-IP**, **X-Authenticated-User**, and **X-Authenticated-Groups** header extensions in an ICAP message to provide information about the source of the encapsulated HTTP message.

  [See ICAP Service Redirect and icap-redirect.]

- **Express Path session status CLI monitoring improvement and traffic logging (SRX4600, SRX5400, SRX5600, and SRX5800)**—The Express Path (formerly known as services offloading) support is already available on SRX4600, SRX5400, SRX5600, and SRX5800 Series devices. Express Path considerably reduces packet-processing latency. Starting in Junos OS Release 20.1R1, you can view the total number of services-offload sessions and total number of services-offload packets processed in the CLI. In addition, you can configure the services-offload traffic logging at the logical system and tenant system level.

  [See Express Path.]

VPNs

- **Common configuration payload password support for RADIUS server (SRX Series and vSRX)**—Starting in Junos OS Release 20.1R1, you can configure a common password for IKEv2 configuration payload requests for an IKE gateway configuration. The common password in the range of 1 to 128 characters allows the administrator to define a common password. This password is used when the SRX Series device is requesting an IP address on behalf of a remote IPsec peer using IKEv2 configuration payload over the RADIUS server. The RADIUS server matches the credentials before it assigns any IP information to the configuration payload request.

  [See Understanding IKEv2 Configuration Payload.]

SEE ALSO

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 20.1R3

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.]

- **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).]

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.]

Routing Protocols

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the lsdist.0 and lsdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, we added multiple secondary loopback addresses in the traffic engineering database to the lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router ID.

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.]

What's Changed in Release 20.1R2

ATP Cloud

- **Dynamic address entries on SRX Series devices in chassis cluster mode**—Starting in Junos OS Release 20.1R2, for SRX Series devices in chassis cluster mode, the dynamic address entry list is retained on the device even after the device is rebooted following a loss of connection to Juniper Advanced Threat Prevention Cloud (ATP Cloud).

Flow-Based and Packet-Based Processing

- On SRX Series devices in earlier releases, when the session table was full there was no alarm set to indicate this. Starting from this release, when the percent of flow session table utilization is 95% on FPC and PIC, an alarm message Flow session table is almost full on FPC <number> PIC <number> is set. Similarly, when the percent of DCP session table utilization is 95% on FPC and PIC, an alarm message DCP session table is almost full on FPC <number> PIC <number> is set.

  [See Understanding Session Cache.]

J-Web

- Change in the J-Web browser tab title (SRX Series)—Starting in Junos OS Release 20.1R2, the J-Web browser tab title displays the device model and the hostname. The same details are displayed when you hover over the J-Web browser tab.

  For example, when you access J-Web for an SRX320 device with a host name srx320-xyz, the J-Web browser tab displays the title as *J-Web (srx320 – srx320-xyz)*.

  If the hostname is not configured, you can see the host URL or IP address in the J-Web browser tab title. For example, *J-Web (srx320 – <device IP address>)*.

VPNs

- **The junos-ike package installed by default (SRX5000 Series devices)**— For SRX5000 Series devices with RE3 installed, the junos-ike package is installed by default. As a result, iked and ikemd process runs on the Routing Engine by default instead of IPsec key management daemon (kmd). In earlier Junos OS Releases, junos-ike package is an optional package for SRX5000 Series devices with RE3 and IPsec Key Management Daemon (KMD) runs by default.

  [See Enabling IPsec VPN Feature Set on SRX5K-SPC3 Services Processing Card.]

- **IKE Index displayed in show security ipsec security-associations detail Output (SRX5400,SRX5600, SRX5800)**— When you execute the **show security ipsec security-associations detail** command, a new output field **IKE SA Index** corresponding to every IPsec Security Association (SA) within a tunnel is displayed under each IPsec SA information.

  [See show security ipsec security-associations.]

What's Changed in Release 20.1R1

ALG

- **Disable the do not fragment flag from packet IP header (SRX Series and vSRX)**—Starting in Junos OS Release 20.1R1, we've introduced the **clear-dont-frag-bit** option at the [**edit security alg alg-manager**] hierarchy level to disable the do not fragment flag from the packet IP header, which allows the packet to be split after NAT is performed.

  In Junos OS releases earlier than Release 20.1R1, when the ALG performs payload-NAT, sometimes the size of the packet becomes bigger than the outgoing interface maximum transmission unit (MTU). If the packet IP header has the do not fragment flag, this packet cannot be sent out.

  [See alg-manager.]

Application Security

- Starting in Junos OS Release 20.1R1, you can enable application identification (AppID) to classify a web application that is hosted on a content delivery network (CDN) such as AWS, Akamai, Azure, Fastly, and Cloudflare and so on accurately. Use the following configuration statement to enable CDN application classification:

  ```
  [edit]
  user@host# user@hots# set service application-identification enable-cdn-application-detection
  ```

  When you apply the configuration, AppID identifies and classifies actual applications that are hosted on the CDN.

  [See Application Identification]

- You can configure maximum memory limit for the deep packet inspection (DPI) by using the following configuration statement:

  ```
  user@host# set services application-identification max-memory memory-value
  ```

  You can set 1 through 200000 MB as memory value.

  Once the JDPI memory consumption reaches to 90% of the configured value, then DPI stops processing new sessions.

  [See Application Identification]

- Starting in Junos OS Release 20.1R1, you can configure and use IP protocol-based custom application signatures on your SRX Series device. In previous versions of Junos OS Releases from 19.2 through 19.4 release, IP protocol based custom application signatures did not work as expected.

  In Junos OS Releases in 19.2 through Junos OS Releases 19.4 and their maintenance releases, IP protocol based custom application signatures do not work as expected. As a workaround, you can configure the IP protocol-based applications at the following hierarchy levels:

- For unified policy: Use service based application configuration as below:

```
user@host# set applications application application-name protocol IP -proto-number
```

- For legacy application firewall: Use predefined IP protocol applications as below:

```
user@host# set security application-firewall rule-sets rule-set-name rule rule-name match dynamic-application
    junos:IPP-IGMP
```

[See Custom Application Signatures for Application Identification.]

Ethernet Switching and Bridging

- **LLDP support on redundant Ethernet interfaces (SRX300, SRX320, SRX340, SRX345, SRX550M, and SRX1500)**—Starting in Junos OS Release 20.1R1, you can configure the Link Layer Discovery Protocol (LLDP) on redundant Ethernet (reth) interfaces. Use the **set protocol lldp interface *<reth-interface>*** command to configure LLDP on the reth interface.

  [See Configuring LLDP and Ethernet Ports Switching Overview for Security Devices.]

J-Web

- Deactivated policy rules are not visible in the J-Web UI (SRX Series)—J-Web does not support disabling or enabling the security firewall or global policy rules from Junos OS Release 20.1R1. The policy rules that are deactivated through CLI are also not visible in the J-Web UI. As a workaround, use CLI to disable or enable the policy rules on the device.

Unified Threat Management (UTM)

- **Increase in the UTM scale number (SRX1500, SRX4100, SRX4200, SRX4600, SRX4800, SRX5400, SRX5600, and SRX5800)**—Starting with Junos OS Release 20.1R1, on SRX Series devices, UTM policies, profiles, MIME patterns, filename extensions, protocol commands, and custom messages are increased up to 1500. Custom URL patterns and custom URL categories are increased up to 3000.

  [See Unified Threat Management overview.]

VPNs

- **Public key infrastructure warning message (SRX Series)**—Starting in Junos OS Release 20.1R1, a warning message **ECDSA Keypair not supported with SCEP for cert_id *<certificate id>*** is displayed when you try to enroll a local certificate using an Elliptic Curve Digital Signature Algorithm (ECDSA) key with Simple Certificate Enrollment Protocol (SCEP) because ECDSA key is not supported with SCEP.

  Prior to Junos OS Release 20.1R1, the warning message is not displayed.

  [See Example: Enrolling a Local Certificate Online Using SCEP.]

- **Change in display of local certificate serial number (SRX Series)**—In Junos OS Release 20.1R1, the output of the **show security pki local-certificate detail** command is modified to display the PKI local certificate serial number with **0x** as prefix to indicate that the PKI local certificate is in the hexadecimal format.

  [See show security pki local-certificate (View).]

SEE ALSO

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for SRX Series.

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

General Routing

- On an SRX4600 device, when LLDP is configured on the interfaces, Packet Forwarding Engine stops operating is seen due to the segmentation problem. LLDP is not supported on SRX4600 currently, but can be configured. PR1422466
- On SRX5400, SRX5600, and SRX5800 devices, on reth interfaces that are configured as DHCP clients, after a reboot of the device the interface might not get an IP address when you use the default number

of DHCP retransmission attempts. When the number of retransmission attempts is increased to 5 or higher, it works fine. PR1458490

- MACsec is not working as expected on ports of the SRX380 device with peer interfaces on the same cluster. PR1479705
- 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

J-Web

- When a dynamic application is created for an edited policy rule, the list of services is blank when the Services tab is clicked and then the policy grid is autorefreshed. As a workaround, create a dynamic application as the last action while modifying the policy rule and click the Save button to avoid loss of configuration changes made to the policy rule. PR1460214

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
- On the SRX5000 line of devices with SPC3 cards, sometimes the IKE SA is not seen on the device when an st0 binding on a VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). PR1441411
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, with 60,000 tunnels up, when RG0 failover happens while an IPsec and/or IKE rekey is in progress, those rekeying tunnels might go down and traffic loss might be seen until the tunnel is reestablished. PR1471499
- 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
- After IPsec tunnel using policy-based VPN is overwritten by another VPN client, traffic using this IPsec tunnel will be dropped. PR1546537

SEE ALSO

What's New | 312

What's Changed | 319

Open Issues

IN THIS SECTION

Learn about open issues in this release for SRX Series.

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 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

- Command show security pki local-certificate logical-system all is not showing any output. PR1414628
- On an SRX340 device with J-Flow version 9 configured, the flowd process might generate core files frequently when the device is busy. PR1463689
- The firewall Web authentication graphics have been updated. PR1482433

- CLI autocomplete is now available for both secintel and advanced anti-malware products. PR1487419
- The SRX380 device has a large number of switch ports and needs to monitor their statistics in real time. Due to this, some of the Routing Engine CPU is always consumed during these operations. If you need to ensure maximum Routing Engine performance, you can dedicate a Packet Forwarding Engine core to the uKernel to more efficiently manage this task. To do this, use the set chassis dedicated-kern-cpu command. PR1527147
- 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

Intrusion Detection and Prevention (IDP)

- Starting from Junos OS Release 21.1, either greater-than or less-than 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. Anomalies and generic attacks will be part of all groups created. PR1397599

J-Web

- On the SRX5000 line of devices, J-Web might not be responsive sometimes when you commit configuration changes after adding a new dynamic application while creating a new firewall rule. J-Web displays a warning while validating the configuration due to dynamic application or any other configuration changes. As a workaround, refresh the J-Web page. PR1460001

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

VPNs

- 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

- On the SRX5000 line of devices with SPC3 cards, sometimes IKE SA is not seen on the device when the st0 binding on the VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). PR1441411
- In an IPsec VPN scenario on SRX5400, SRX5600, and SRX5800 platforms, the iked process treats retransmission of IKE_INIT request packets as new connections when the SRX Series device acts as a responder of IKE negotiation. This causes IKE tunnel negotiation to fail, and IPsec VPN traffic might be impacted. PR1460907
- On the SRX5000 line of devices with SPC3 and SPC2 mixed mode, with a very large number of IKE peers (60,000) with dead peer detection (DPD) enabled, IPsec tunnels might flap in some cases when IKE and IPsec rekeys are happening at the same time. PR1473523
- Some TCP connections going through IPsec tunnels are getting stuck after RG1 failover. PR1477184
- During 10,000 tunnel ramp-up, sometimes, IKED generates a core file. PR1479548
- Unexpected extra characters NL were seen with PyEZ XML outputs. This caused issues while writing op-scripts. However, with normalize=True in PyEz script, you can avoid having NL between each tag and with the pretty_print option ensure that the print is clean. PR1492146
- The SRX5000 line of devices with SPC3 does not support simultaneous IKE negotiation in Junos OS Releases 19.2, 19.3, 19.4, and 20.1. PR1497297

SEE ALSO

What's New | **312**

What's Changed | **319**

Known Limitations | **325**

Resolved Issues | **329**

Documentation Updates | **343**

Migration, Upgrade, and Downgrade Instructions | **343**

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: 20.1R3

Chassis Clustering

- Disabled node on SRX chassis cluster sends out ARP request packets. PR1548173
- SPU might stop under GPRS tunneling protocol scenario. PR1559802

Flow-Based and Packet-Based Processing

- The rst-invalidate-session command does not work if configured together with the no-sequence-check command. PR1541954

Forwarding and Sampling

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

General Routing

- During an upgrade, system displays the following incorrect license warnings when utilizing licensable features even if the license is present on the device: requires 'idp-sig' license. PR1519672
- Packet drops might be seen with all commit events with 1G speed configured interface. PR1524614
- The MAC table is null in Layer 2 mode after one pass-through session is created successfully. PR1528286
- Junos OS: Memory leak when querying Aggregated Ethernet (AE) interface statistics (CVE-2021-0230). PR1528605
- 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
- The JNH memory might leak on the Trio-based line cards. PR1542882
- Tail drops might occur on SRX Series devices if shaping-rate is configured on lt interface. PR1542931
- The kmd process might crash when the interface flaps. PR1544800
- The flowd process might crash on SRX Series devices. PR1545628
- SRX1500 reports fans running at over speed. PR1546132
- On SRX4100 and SRX4200, if PEM0 is removed, the output of jnxOperatingDescr.2 might be incomplete. PR1547053
- Advanced anti-malware file or email statistics does not get incremented with the latest PB version. PR1547094

- On vSRX2.0, vSRX3.0, SRX1500, SRX4100, SRX4200, SRX4600 running chassis cluster in Junos OS Release 18.3 or later releases, multiple messages of "LCC: ch_cluster_lcc_set_context:564: failed to lock chassis_vmx mutex 11" are generated in the chassisd log file. These messages may recur after every few seconds and they do not have any impact on system operation. PR1547953
- When Junos OS software is upgraded to Junos OS Release 20.3, you might see the error "ERROR: Failed to setup symlinks in alternate root". PR1548626
- PKI CMPv2 client certificate enrolment does not work on SRX Series devices when using root-CA. PR1549954
- Lcmd log "gw_cb_presence:136: PEM(slot = 0): error detecting presence (fruid = 15, drv_id = 30, status = -11)" generates every second on the SRX4100 and SRX4200 devices. PR1550249
- The speed mismatch error is seen while trying to commit reth0 with gigether-options. PR1553888
- On the SRX550M device, the dumpdisklabel command fails with message "ERROR: Unknown platform srx550m". PR1557311
- The idpd process might stop when committing IDP configuration under logical systems and tenant systems during RGs failover. PR1561298
- The flowd process might stop and generates a core files if Jflow V9 is configured. PR1567871
- Wi-Fi mPIM on SRX Series devices is reaching out to NTP and DNS servers. PR1569680
- MACsec not using network-control queue. PR1571977
- Traffic going through the VRRP interface might be dropped when VRRP enabled IRB interface goes down. PR1572920
- The 1G interfaces might not come up after device reboot. PR1585698

Interfaces and Chassis

- When SRX Series devices receive proxy ARP requests on VRRP interfaces, SRX Series devices send ARP replies with the underlying interface MAC address. PR1526851
- Backup Routing Engine or backup node may stuck in bad status with improper "backup-router" configuration. PR1530935

Intrusion Detection and Prevention (IDP)

- The flowd or srxpfe process might generate core files during the idpd process commit on SRX Series devices. PR1521682
- Need system log to indicate signature download completion. PR1543571
- IDP policy load might fail post image upgrade for Junos OS 15.1x49 releases. PR1546542
- The idpd process might stop and generates a core files. PR1547610
- The IDP policy process might become unresponsive and fail to compile the IDP policy after an IDP automatic update. PR1577684

J-Web

- The "+" button is not shown in the J-Web interface menu. PR1550755

Network Management and Monitoring

- The mib2d process pause and generates a core file on backup Routing Engine. PR1557384

Platform and Infrastructure

- Continuous L2ALD and L2ALM log messages seen on nodes of chassis cluster of SRX5000 line of devices. PR1501752
- 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 the reth interface is deleted. PR1538273

Routing Policy and Firewall Filters

- The flowd or srxpfe process might stop when an SRX Series or NFX Series device running Junos OS Release 18.2R1 or later supports the unified policy feature. PR1544554
- 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 PFE CPU utilization. PR1549366
- Policy configured with "route-active-on" condition may incorrectly work for local routes. PR1549592
- 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
- On the SRX5000 line of devices, the secondary node might get stuck in performing ColdSync after a reboot, upgrade, or if ISSU is performed. PR1558382
- Traffic loss might be seen when a big number of applications or addresses is referenced by one policy. PR1576038

Unified Threat Management (UTM)

- Stream buffer memory leak might happen when UTM is configured under unified policies. PR1557278
- UTM license expiry event loss may cause the device to not quit the advanced service mode and maximum-sessions is decreased by half. PR1563874

VPNs

- IPsec tunnel could flap when ESN is enabled. PR1488087
- IPsec SA is missing the keyword NULL after RG failover. PR1507270
- The traffic might be dropped when IPsec VPN with NAT-T enabled. PR1522017
- On all SRX series devices using IPsec with NAT Traversal, MTU size for the external interface might be changed after IPsec SA is re-established. PR1530684

- The flowd process might stop during IPsec SA renegotiation on SRX5000 line of devices. PR1545916
- After the IPsec tunnel using policy-based VPN is overwritten by another VPN client, traffic using this IPsec tunnel will be dropped. PR1546537
- Traffic going through policy-based IPsec tunnel might be dropped after RG0 failover. PR1550232
- A session might be closed when the session is created during the IPsec rekey. PR1564444
- When there are multiple IPsec SA, backup SA start IPsec rekey. PR1565132
- SPI mismatch caused by simultaneous rekeys under kmd stress. PR1571105

Resolved Issues: 20.1R2

Application Layer Gateways (ALGs)

- FTPS traffic might get dropped on SRX Series and MX Series devices if FTP ALG is used. PR1483834
- The srxpfe and mspmand process might stop if FTPS is enabled in a specific scenario. PR1510678

Application Security

- AppQoE support for dynamic-application. PR1503400

Chassis Clustering

- The show chassis temperature-thresholds command displays extensive FPC 0 output. PR1485224
- On SRX4100 and SRX4200 devices with chassis cluster in transparent mode, when a failover occurs for RG1, the interface on the new secondary node is getting flapped as expected to let the switch update its MAC address table. PR1490291
- IP monitoring on SRX4100 and SRX4200 device might fail in the rare event that a chassis internal connection between Routing Engine and Packet Forwarding Engine is temporarily down after RG0 failover. PR1502462
- The ISSU fails with timeout due to cold synchronization failure. PR1502872

Flow-Based and Packet-Based Processing

- The show security group-vpn server statistics |display XML command output is not in expected format. PR1349959
- ECMP load balancing does not happen when RG1 node 0 is secondary. PR1475853
- The flowd or srxpfe process might stop when deleting the user firewall local authentication table entry. PR1477627
- On Web proxy, memory leak in association hash table and DNS hash table. PR1480760
- IMAP curl sessions stuck in the active state if AAMW IMAP block mode is configured. PR1484692
- The flowd process might stop and impact services if J-Flow version 9 is configured. PR1486528

- The configuration set chassis psu redundancy n-plus-n needs support on in high availability (HA) mode. PR1486746
- Commit does not work after the installation through boot loader. PR1487831
- If a cluster ID of 16 or multiples of 16 is used, the chassis cluster might not come up. PR1487951
- CPU board inlet increases after OS upgrade from Junos OS Release 15.1X49 to Junos OS Release 18.x. PR1488203
- All interfaces remain in the down status after the SRX300 line of devices power up or reboot. PR1488348
- GRE or IPsec tunnel might not come up when the set security flow no-local-favor-ecmp command is run. PR1489276
- Sometimes multiple flowd core files are generated on both nodes of a chassis cluster at the same time when changing media MTU. PR1489494
- Continuous drops seen in control traffic, with high data queues in one SPC2 PIC. PR1490216
- Not able to clear the warm sessions on the peer SRX Series devices. PR1493174
- Phone client stop seen during SRX345 device ZTP with CSO. PR1496650
- Outbound SSH connection flap or memory leak issue might be observed during the high rate of pushing the configuration to the ephemeral database. PR1497575
- Unexpected flow logging traffic beyond the packet filter. PR1497939
- Traffic interruption happens due to MAC address duplication between two devices running Junos OS. PR1497956
- Don't use capital characters for source-identity when using the show security match-policies command. PR1499090
- On SRX Series devices, when the GRE or IP-IP tunnel is used, if some interface change events happen (such as, interface flapping), traffic drop might be seen. PR1500091
- The srxpfe or flowd process might stop due to memory corruption within JDPI. PR1500938
- J-Flow version 9 does not display the correct outgoing interface for APBR traffic. PR1502432
- A condition within TCP proxy could result in downloads becoming permanently stuck or not completing. TCP proxy is used by multiple services, including Juniper ATP Cloud in block mode, ICAP, SSL proxy, antivirus, content filtering, and antispam. PR1502977
- Fabric interface might be monitored down after chassis cluster reboot. PR1503075
- A cfmd core file is observed when LTM is triggered for the session configured on the ethernet-switching interface without bridge domain configuration. PR1503696
- Layer 2 ping is not working with remote MEP. PR1504986
- SOF asymmetric scenario is not working with the phase 1 solution. PR1507865

- If the dynamic-app configured along with other Layer 7 application in different rules, after App identified still the SLA database with application any showing up some sessions. PR1514973
- VRRP does not work on the redundant Ethernet interface with a VLAN ID greater than 1023. PR1515046
- A logic issue was corrected in SSL proxy that could lead to an srxpfe or flowd core file under load. PR1516903
- The PPPoE session does not come up after return to zero on SRX Series devices. PR1518709
- The TCP packet might be dropped if syn-proxy protection is enabled. PR1521325
- On SRX Series devices with chassis clusters, high CPU usage might be seen due to the llmd process. PR1521794
- Certificate validation might fail when OCSP is used and the OCSP server is a dual-stack device. PR1525924
- Traffic rate shown in the CLI command is not accurate. PR1527511
- On SRX4100 and SRX4200 devices, four out of eight fans might not work. PR1534706
- The rst-invalidate-session configuration does not work if configured together with no-sequence-check. PR1541954
- NSD core file is generated at function nsd_malloc, file ../../../../../../src/usp/usr.sbin/nsd/common/nsd_common.c, line 482. PR1542942

Interfaces and Chassis

- PPO IPv6 route does not work. PR1495839

Intrusion Detection and Prevention (IDP)

- When intelligent inspection status changes, syslog is not getting generated on SRX300 and SRX500 lines of devices. PR1448365
- The IDP attack detection might not work in a specific situation. PR1497340
- IDP's custom-attack time-binding interval command was mistakenly hidden within the CLI. PR1506765

J-Web

- While creating a firewall policy rule, the list of available dynamic applications is empty in HA on the Select Dynamic Application page. PR1490346
- Junos OS: Reflected Cross-site Scripting vulnerability in J-Web and web based (HTTP/HTTPS) services (CVE-2020-1673) PR1493385
- You cannot configure Redundant PSU and Power Budget Statistics on the SRX380 device, which is in HA mode, through J-Web. PR1493713
- The J-Web users might not be able to configure PPPoE using the PPPoE wizard. PR1502657
- J-Web chassis status widget is incorrectly reporting temperature alarms. PR1507156
- The parameters show another LSYS at J-Web in a multiple logical systems scenario. PR1518675

Layer 2 Ethernet Services

- DHCP does not work after running request system zeroize or load factory-default. PR1521704

MPLS

- BGP session might keep flapping between two directly connected BGP peers because of the wrong TCP-MSS in use. PR1493431

Network Address Translation (NAT)

- Not all NAT sessions are synchronized from Node 1 to Node 2. PR1473788
- Issuing the show security nat source paired-address command might return an error. PR1479824

Platform and Infrastructure

- On the SRX1500 device and the SRX4000 line of devices, physically disconnecting the cable from the fxp0 interface causes hardware monitor failure and redundancy group failover, when the device is the primary node in a chassis cluster. PR1467376
- The SRX1500 device and the SRX4000 line of devices might boot up with the rescue configuration after a power outage. PR1490181
- Packets get dropped when the next hop is IRB over the LT interface. PR1494594
- The /usr/libexec/ui/yang-pkg and /usr/libexec/ui/pyang files are not found in SRX Series devices during YANG installation. PR1496577
- Junos OS: Arbitrary code execution vulnerability in Telnet server (CVE-2020-10188). PR1502386
- On the SRX1500 device, the factory-default configuration for ge-0/0/0 and ge-0/0/15 should be set with family inet DHCP. PR1503636
- Syslog reporting "PFE_FLOWD_SELFPING_PACKET_LOSS: Traffic impact: Selfping packets loss/err: 300 within 600 second" error messages in node 0 and node1 control panel. PR1522130

Routing Policy and Firewall Filters

- TCP proxy was mistakenly engaged in unified policies when Web filtering was configured in potential match policies. PR1492436
- Traffic might fail to hit policies if match dynamic-application and match source-end-user-profile options are configured under the same security policy name. PR1505002

Routing Protocols

- The BGP route-target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. PR1492743
- The rpd might report 100% CPU usage with BGP route damping enabled. PR1514635

Unified Threat Management (UTM)

- UTM websense redirect supports IPv6 messages. PR1481290

- UTM doesn't let e-mails from outside to inside to be received. PR1523222

VPNs

- IKE SA does not get cleared and is showing very long lifetime. PR1439338
- With NCP remote access solution, in a PathFinder case (for example, where IPsec traffic has to be encapsulated as TCP packets), TCP encapsulation for transit traffic is failing. PR1442145
- The newly configured IPsec tunnels might be stuck in VPNM verify-path state in a tunnel scaled scenario. PR1464353
- On an SRX4200 device, 35 percent of drop is seen in all TPS cases. PR1481625
- On SRX Series devices with SPC3, when overlapping traffic selectors are configured, multiple IPsec SAs get negotiated with the peer device. PR1482446
- Traffic might be lost after the rekey if SRX Series devices responder-only is configured. PR1485029
- Use different XML tags for local and remote IKE IDs to avoid confusion. PR1493368
- Issue with XML RPC show security ipsec tunnel-distribution summary output. PR1494274
- On SRX Series devices using IPsec with NAT traversal, MTU size for the external interface might be changed after IPsec SA is reestablished. PR1530684

Resolved Issues: 20.1R1

Application Layer Gateways (ALGs)

- Packet's IP header have DF flag might be dropped by SRX Series ALG after payload-NAT. PR1444068
- On the SRX5000 line of devices, the H323 call with NAT64 could not be established. PR1462984
- RTSP data sessions are cleared unexpectedly during cold sync. PR1468001
- The flowd or srxpfe process might stop when an ALG creates a gate with an incorrect protocol value. PR1474942
- SIP messages that need to be fragmented might be dropped by SIP ALG. PR1475031

Authentication and Access Control

- Same-source IP sessions are cleared when the IP entry is removed from the UAC table. PR1457570

Chassis Clustering

- IP monitoring might fail on the secondary node. PR1468441
- An unhealthy node might become primary in SRX4600 devices with chassis cluster scenario. PR1474233

Flow-Based and Packet-Based Processing

- The **trusted-ca** and **root-ca** names or IDs should not be the same within an SSL proxy configuration. PR1420859
- Packet loss is caused by FPGA back pressure on SPC3. PR1429899
- Control logical interface is not created by default for LLDP. PR1436327
- Security logs cannot be sent to the external syslog server through TCP. PR1438834
- The SPC card might stop on the SRX5000 line of devices. PR1439744
- Flowd process core files are generated in the device while testing NAT PBA in AA mode. PR1443148
- The SSL-based AppID simplification effort (removal of HTTPS, POP3S, IMAPS, SMTPS). PR1444767
- In the BERT test for E1 interface, bits counts number is not within the range. PR1445041
- The flowd process might stop on SRX Series devices when chassis cluster and IRB interface are configured. PR1446833
- The AAWM policy rules for IMAP traffic sometimes might not get applied when passed through SRX Series devices. PR1450904
- Introduction of default inspection limits for application identification to optimize CPU usage and improve resistance to evasive applications. PR1454180
- The SRX Series devices stop and generate several core files. PR1455169
- When you try to reset the system configuration on an SRX1500 device using the reset config button, it does not work properly. PR1458323
- The security flow traceoptions fills in with RTSP ALG-related information. PR1458578
- Optimizations were made to improve the connections-per-second performance of SPC3. PR1458727
- LTE dual CPE support with mPIMs when modem receives disconnect event from ISP, need to increase wait timer. PR1460102
- The **security-intelligence** CC feed does not block HTTPS traffic based on SNI. PR1460384
- The AAMWD process exceeds 85 percent RLIMIT_DATA limitation due to memory leak. PR1460619
- Added command to clear specified associated client. PR1461577
- The srxpfe or flowd process might stop if the sampling configuration is changed. PR1462610
- The tunnel packets might be dropped because the gr0.0 or st0.0 interface is wrongly calculated after a GRE or VPN route change. PR1462825
- Fragmented traffic might get looped between the fab interface in a rare case. PR1465100
- TCP session might not time out properly upon receiving TCP RESET packet. PR1467654
- A core file might be generated when you perform an ISSU on SRX Series devices. PR1463159
- The PKI daemon keeps leaking memory on SRX Series devices. PR1465614

- HTTP block message stops working after SNI check for HTTPS session. PR1465626
- Loading CA certificate causes PKI daemon core file to be generated. PR1465966
- The jbuf process usage might increase up to 99 percent after Junos OS upgrade. PR1467351
- The rpd process might stop after several changes to the **flow-spec** routes. PR1467838
- Packet Forwarding Engine might generate core files because SSL proxy is enabled on NFX Series and SRX Series devices. PR1467856
- Server unreachable is detected; ensure that port 443 is reachable. PR1468114
- Tail drop on all ports is observed when any switch-side egress port gets congested. PR1468430
- FTP data connection might be dropped if SRX Series devices send the FTP connection traffic through the dl interface. PR1468570
- RPM test probe fails to show that round-trip time has been exceeded. PR1471606
- Look up failure for expected e-mail address in DUT. PR1472748
- Stateful firewall rule configuration deletion might lead to memory leak. PR1475220
- The **dfs-off** function is enabled. PR1475294
- The nsd process pause might be seen during device reboots if dynamic application groups are configured in policy. PR1478608
- The **show mape rule statistics** command might display negative values. PR1479165
- Sometimes multiple flowd core files are generated on both nodes of chassis cluster at the same time when changing media MTU. PR1489494

Interfaces and Chassis

- The number of mgd processes increases because the mgd processes are not closed properly. PR1439440
- Static route through dl0.0 interface is not active. PR1465199
- MAC limiting on Layer 3 routing interfaces does not work. PR1465366

Intrusion Detection and Prevention (IDP)

- SNMP queries might cause **commit** or **show** command to fail due to IDP PR1444043
- Updating the IDP security package offline might fail in SRX Series devices. PR1466283

J-Web

- The default log query time in J-Web monitoring functionality has been reduced. This increases the responsiveness of the landing pages. PR1423864
- Editing destination NAT rule in J-Web introduces a nonconfigured routing instance field. PR1461599

- The Go button within the J-Web Monitor>Events view now correctly refreshes the logs even when using a blank search query. PR1464593
- J-Web security resources dashboard widget was not being populated correctly. PR1464769

Layer 2 Ethernet Services

- The metric is not changing when configured under the DHCP. PR1461571

Network Address Translation (NAT)

- The flowd or srxpfe process might stop when traffic is processed by both ALGs and NAT. PR1471932
- Issuing the **show security nat source paired-address** command might return an error. PR1479824

Network Management and Monitoring

- The flowd or srxpfe process might stop immediately after committing the jflowv9 configuration or after upgrading to affected releases. PR1471524
- SNMP trap coldStart agent-address becomes 0.0.0.0. PR1473288

Platform and Infrastructure

- Modifying the REST configuration might cause the system to become unresponsive. PR1461021
- VM core files might be generated if the configured sampling rate is more than 65,535. PR1461487
- On the SRX300 line of devices, you might encounter Authentication-Table loading slowly while using user-identification. PR1462922
- The AE interface cannot be configured on an SRX4600 device. PR1465159
- On SRX Series devices, Packet Forwarding Engine memory might be used up if the security intelligence feature is configured. PR1472926
- Support LLDP protocol on reth interface. PR1473456
- Certificate error while configuration validation during Junos OS upgrade. PR1474225
- Packet drop might be observed on the SRX300 line of devices when adding or removing an interface from MACsec. PR1474674
- The commands **request system power-off** and **request system halt** might not work correctly. PR1474985
- The flowd process core files might be seen when there are mixed NAT-T traffic or non-NAT-T traffic with PMI enabled. PR1478812
- When SRX5K-SPC3s or MX-SPC3s are installed in slots 0 or 1 in SRX5800 or MX960 devices, EMI radiated emissions are observed to be higher than regulatory compliance requirements. PR1479001
- The RGx might fail over after RG0 failover in a rare case. PR1479255
- The wl- interface stays in ready status after you execute **request chassis fpc restart** command in Layer 2 mode. PR1479396

- Recent changes to JDPI's classification mechanism caused a considerable performance regression (more than 30 percent). PR1479684
- The flowd or srxpfe process might crash when advanced anti-malware services are used. PR1480005

Routing Policy and Firewall Filters

- Security policies cannot synchronize between Routing Engine and Packet Forwarding Engine on SRX Series devices. PR1453852
- Traffic log shows wrong custom-application name when the **alg ignore** option is used in application configuration. PR1457029
- The NSD process might get stuck and cause problems. PR1458639
- Some domains are not resolved by the SRX Series devices when using DNS address book. PR1471408
- The count option in security policy does not take effect even if the policy count is enabled. PR1471621
- Support for dynamic tunnels on SRX Series devices was mistakenly removed. PR1476530

Routing Protocols

- SSH login might fail if a user account exists in both local database and RADIUS or TACACS+. PR1454177
- The rpd might stop when both instance-import and instance-export policies contain as-path-prepend action. PR1471968

Unified Threat Management (UTM)

- Increase the scale number of UTM profile or policy for the SRX1500 device, and the SRX4000 and SRX5000 lines of devices. PR1455321
- The utmd process might pause after deactivating UTM configuration with predefined category upgrading used. PR1478825

VLAN Infrastructure

- ISSU failed from Junos OS Release 18.4R2.7 to Junos OS Release 19.4, with secondary node PICs in present state after upgrading to Junos OS Release 19.4. PR1468609

VPNs

- IPsec SA inconsistent on SPCs of node0 and node1 in SRX Series devices with chassis cluster. PR1351646
- After RG1 failover, IKE phase 1 SA is getting cleared. PR1352457
- IPsec VPN missing half of the IKE SA and IPsec SA showing incorrect port number when scaling to 1000 IKEv1 AutoVPN tunnels. PR1399147
- The IKE and IPsec configuration under groups is not supported. PR1405840
- The established tunnels might remain unchanged when an IKE gateway is changed from AutoVPN to Site-to-Site VPN. PR1413619

- The VPN tunnel might flap when IKE and IPsec rekey happen simultaneously. PR1421905
- Old tunnel entries might be observed in the output of show security IPsec or IKE SA. PR1423821
- The **show security ipsec statistics** command output displays buffer overflow and wraps around 4,---,---,--- count. PR1424558
- Tunnel does not come up after changing configurations from IPv4 to IPv6 tunnels in the script with gateway lookup failed error. PR1431265
- P1 configuration delete message is not sent on loading baseline configuration if there has been a prior change in VPN configuration. PR1432434
- After a long time (a few hours) of traffic during a mini PDT test, the number of IPsec tunnels is much higher than expected. PR1449296
- Some IPsec tunnels flap after RGs failover on the SRX5000 line of devices. PR1450217
- IPsec VPN flaps if more than 500 IPsec VPN tunnels are connected for the first time. PR1455951
- Traffic is not sent out through an IPsec VPN after update to Junos OS Release 18.2 or later. PR1461793
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, with IKEv1 enabled IKE, the daemon might generate a core file, when IKESA is expired and IPsec tunnel associated with the expired IKESA exists in case of an RG0 failover. Daemon recovers eventually. PR1463501
- The IPsec VPN tunnels cannot be established if overlapped subnets are configured in traffic selectors. PR1463880
- IPsec tunnels might lose connectivity on SRX Series devices after chassis cluster failover when using AutoVPN point-to-multipoint mode. PR1469172
- IPsec tunnels might flap when one secondary node is coming online after reboot in SRX Series high availability environment. PR1471243
- The kmd process might crash continually after the chassis cluster failover in the IPsec ADVPN scenario. PR1479738

SEE ALSO

What's New | **312**

What's Changed | **319**

Known Limitations | **325**

Open Issues | **327**

Documentation Updates | **343**

Migration, Upgrade, and Downgrade Instructions | **343**

Documentation Updates

IN THIS SECTION

Dynamic Host Configuration Protocol (DHCP)

- **Introducing DHCP User Guide**—Starting in Junos OS Release 20.1R1, we are introducing the DHCP User Guide for Junos OS routing, switching, and security platforms. This guide provides basic configuration details for your Junos OS device as DHCP Server, DHCP client, and DHCP relay agent.

  [See DHCP User Guide.]

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

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

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

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

provide direct upgrade and downgrade paths. You can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 15.1X49, 17.3, 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from one Junos OS Release to the next release or one release after the next release. For example you can upgrade from Junos OS Release 15.1X49 to Release 17.3 or 17.4, Junos OS Release 17.4 to Release 18.1 or 18.2, and from Junos OS Release 18.1 to Release 18.2 or 18.3 and so on.

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

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

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

For information about ISSU, see the Chassis Cluster User Guide for Security Devices.

SEE ALSO

- What's New | 312
- What's Changed | 319
- Known Limitations | 325
- Open Issues | 327
- Resolved Issues | 329
- Documentation Updates | 343

Junos OS Release Notes for vMX

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 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.

Release 20.1R3 New and Changed Features

There are no new features for vMX in Junos OS Release 20.1R3.

Release 20.1R2 New and Changed Features

There are no new features for vMX in Junos OS Release 20.1R2.

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 20.1R3

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.]

What's Changed in Release 20.1R2

There are no changes in behavior or syntax for vMX in Junos OS Release 20.1R2.

Known Limitations

There are no known behaviors and limitations for vMX in Junos OS Release 20.1R3.

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

- On a vMX instance, the performance of X710 NIC is lower compared to that of 82599 NIC. 40G line rate can be acheived at 512 byte packet size for X710 NIC as compared to 256 bytes for 82599 NIC. PR1281366
- 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

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: 20.1R3

Platform and Infrastructure

- Multiple vmxt processes might generate core files. PR1534641

Resolved Issues: 20.1R2

Platform and Infrastructure

- In vMX instances, after every commit, the following error message is displayed in the log message: chassisd[7836]: %DAEMON-3-CHASSISD_IOCTL_FAILURE: acb_get_fpga_rev: unable to get FPGA revision for Control Board (Inappropriate ioctl for device). PR1477941
- On vMX instances, configuring the ranges statement for autosensed VLANs (either stacked VLANs or single-tagged VLANs) might not work. This is because the VLANs are not programmed on the NIC drivers. PR1503538
- On vMX, core.vmxt.mpc0' seen at 5 0x096327d5 in l2alm_sync_entry_in_pfes (context=0xd92e7b28, sync_info=0xd92e7a78) at ../../../../../src/pfe/common/applications/l2alm/l2alm_common_hw_api.c:1727. PR1430440
- On MX150 and vMX platform, if flow caching is enabled, VXLAN packets might be discarded. This is because flow caching does not support VXLAN. PR1466470

Licensing

Starting in Junos OS Release 19.2R1, Juniper Agile Licensing introduces a new capability that significantly improves the ease of license management network wide. The Juniper Agile License Manager is a software application that runs on your network and provides an on-premise repository of licenses that are dynamically consumed by Juniper Networks devices and applications as required. Integration with Juniper's Entitlement Management System and Portal provides an intuitive extension of the existing user experience that enables you to manage all your licenses.

- The Agile License Manager is a new option that provides more efficient management of licenses, but you can continue to use individual license keys for each device if required.
- To use vMX or vBNG feature licenses in Junos OS Release 19.2R1 version, you need new license keys. Previous license keys will continue to be supported for previous Junos OS releases, but for the Junos OS 19.2R1 Release and later you need to carry out a one-time migration of existing licenses. Contact Customer Care to exchange previous licenses. Note that you can choose to use individual license keys for each device, or to deploy Agile License Manager for more efficient management of licenses.
- For more information about Agile Licensing keys and capabilities, see Juniper Agile Licensing portal FAQ.

See Juniper Agile Licensing Guide for more details on how to obtain, install, and use the License Manager.

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 vSRX

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 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 Release 20.1R3

There are no new features for vSRX in Junos OS Release 20.1R3.

What's New in Release 20.1R2

There are no new features for vSRX in Junos OS Release 20.1R2.

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 20.1R3

Junos OS 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.]

What's Changed in Release 20.1R2

Platform and Infrastructure

- **Repetition of WALinuxAgent logs causing file size increase (vSRX 3.0)**—The Azure WALinuxAgent performs the provisioning job for the vSRX instances. When a new vSRX instance is deployed, the continued increasing size of the waagent log file might cause the vSRX to stop.

  If the vSRX is still operating, then delete the /var/log/waagent.log directly or run the clear log waagent.log all command to clear the log file.

  Or you can run the set groups azure-provision system syslog file waagent.log archive size 1m and set groups azure-provision system syslog file waagent.log archive files 10 commands to prevent the growing of the waagent logs. These configurations will cause the rotation of log of waagent with the size bigger than 1MB and set a maximum of 10 backups.

  See vSRX with Microsoft Azure.

Known Limitations

IN THIS SECTION

Learn about known limitations in Junos OS Release 20.1R3 for vSRX Series. 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

- On vSRX, when using IPsec VPN tunnels, we recommend that you use GCM encryption algorithms, such as aes-128-gcm. GCM encryption algorithms have better performance on vSRX than CBC encryption algorithms, such as aes-128-cbc. PR1444022

General Routing

- On vSRX 3.0 running on Azure, there might be one more IP address 1.1.1.1 configured on fxp0 intermittently besides the IP assigned by DHCP, which would cause CLI upgrade failure when HSM is enabled. PR1461678
- For vSRX3.0 on Azure, when HSM is enabled, do not use underscore "_" in the certificate id field while creating keypairs. This is a limitation from the Azure KeyVault. PR1475254

J-Web

- When a dynamic application is created for an edited policy rule, the list of services is blank when the Services tab is clicked and then the policy grid is autorefreshed. As a workaround, create a dynamic application as the last action while modifying the policy rule and click the Save button to avoid loss of configuration changes made to the policy rule. PR1460214

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Release 20.1R3 for vSRX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online Junos Problem Report Search application.

Intrusion Detection and Prevention (IDP)

- IDP db file format or convention has changed in Junos OS Release 15.1X49 and later releases. So, if the IDP configuration contains some predefined attacks or attack-groups related configurations, then the system will go to amnesiac mode after upgrade. This is due to the failure in IDP configuration commit. PR1455125

Resolved Issues

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: 20.1R3

General Routing

- The control link might be broken when there is excessive traffic load on the control link in vSRX cluster deployment. PR1524243
- Packet drops might be seen with all commit events with 1G speed configured interface. PR1524614

Interfaces and Chassis

- Junos OS: SRX1500, SRX4100, SRX4200, SRX4600, SRX5000 Series with SPC2/SPC3, vSRX Series: In a multi-tenant environment, a tenant host administrator may configure logical firewall isolation affecting other tenant networks (CVE-2021-0235). PR1537491

Intrusion Detection and Prevention (IDP)

- The flowd or srxpfe process might generate core files during the idpd process commit. PR1521682
- Application identification related signatures might not get triggered. PR1588450

VPNs

- The flowd might stop in IPsec VPN scenario. PR1517262
- A session might be closed when the session is created during the IPsec rekey. PR1564444

Resolved Issues: 20.1R2

Application Layer Gateways (ALGs)

- Previously, the MSRPC ALG only supported operation number 4 messages (opnum 4 - RemoteCreateInstance) for extracting for MSRPC data sessions. We now support opnum 3 messages (opnum 3 - RemoteGetClassObject) for extracting for MSRPC data sessions. PR1462692
- FTPS traffic might get dropped on SRX Series or MX Series platforms if FTP ALG is used. PR1483834

Application Security

- When destination-path-group is deleted in the configuration and added again, the fc-id, dscp, fc name, and loss priority fields are reset. PR1489948
- The flow performance might be reduced in the Security Intelligence scenario. PR1491682
- The flowd srxpfe process might stop when SSL proxy and AppSecure process traffic simultaneously. PR1516969

Flow-Based and Packet-Based Processing

- Traffic drop might be seen when GRE or IP-IP tunnel is used. PR1500091
- The control link might be broken when there is excessive traffic load on the control link in vSRX cluster deployment. PR1524243

Intrusion Detection and Prevention (IDP)

- The IDP attack detection may not work in a specific situation. PR1497340

J-Web

- While creating a firewall policy rule, the list of available dynamic applications is empty in HA on the Select Dynamic Application page. PR1490346
- Infinite loading circle may be encountered through J-Web. PR1493601

Platform and Infrastructure

- On Microsoft Azure deployments, SSH public key authentication is not supported for vSRX 3.0 CLI and portal deployment. PR1402028
- On vSRX platforms with the class-of-service (CoS) feature used, in the rare condition of accessing stale CoS-related memory, the srxpfe process might crash. PR1474124
- The vSRX instance might restart unexpectedly. PR1479156
- An ipfd core is file generated at 0x08601e14 in ipid_msg_process (svr=< optimized out>, client_id=< optimized out>, msg=< optimized out>, len=< optimized out>) at ../../../../../../src/pfe-shared/include/jnx/usp/ipid_shared.h:622. PR1482947
- Cache entries are not seen when global ASC is off. PR1483928

- The srxpfe and flowd crash might be observed in both nodes after failover if UAC is configured. PR1491635
- The clock drift issue might cause control link failure of a vSRX cluster running on the KVM hypervisor. PR1496937
- On vSRX the interfaces might remain shut as the FPC faces issues while coming online after an upgrade attempt on the device. PR1499092
- When SSL proxy is enabled and if the vSRX instance runs out of memory, then the SSL proxy module might stop. PR1505013
- Changes to the configuration command for assigning more vCPUs to the Routing Engine. PR1505724
- In vSRX 3.0 on Azure with keyvault enabled, change in MEK results in deletion of certificates. PR1513456
- With CSO SD-WAN configuration loaded, the flowd process generates core files while deleting the GRE IPsec configuration. PR1513461

Routing Policy and Firewall Filters

- Traffic might fail to hit policies if match dynamic-application and match source-end-user-profile options are configured under the same security policy name. PR1505002

Unified Threat Management (UTM)

- The source and destination IP or port fields were reversed for Content-Filtering and Anti-Virus logs. These fields now reflect the source and destination of the flow correctly. PR1499327

VPNs

- The Ping-icmp test fails after configuring ECMP routes over multipoint tunnel interface VPNs. PR1438311
- The flowd process might stop in an IPsec VPN scenario. PR1517262

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 20.1R3 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, or 19.2 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 20.1R3 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
       20.4K 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 20.1R3 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-x86-64-20.1R3-2021-5-10.0_RELEASE_20.1R3_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-x86-64-20.1R3-2021-5-10.0_RELEASE_20.1R3_THROTTLE.tgz
 no-copy no-validate reboot
Verified junos-vsrx-x86-64-20.1R3-2021-5-10.0_RELEASE_20.1R3_THROTTLE signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING:     This package will load JUNOS 20.1R3 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-20.1R3-2021-5-10.0_RELEASE_20.1R3_THROTTLE.tgz
Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31
junos-srx-mr-vsrx-20.1R3-2021-5-10.0_RELEASE_20.1R3_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31
junos-srx-mr-vsrx-20.1R3-2021-5-10.0_RELEASE_20.1R3_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-20.1R3-2021-5-10.0_RELEASE_20.1R3_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-20.1R3-2021-5-10.0_RELEASE_20.1R3_THROTTLE-linux.tgz
...
```

```
upgrade_platform: Input package
/var/tmp/junos-srx-mr-vsrx-20.1R3-2021-5-10.0_RELEASE_20.1R3_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-20.1R3-2021-5-10.0_RELEASE_20.1R3_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-20.1R3-2021-5-10.0_RELEASE_20.1R3_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 20.1R3 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 20.1R3-2021-5-10.0_RELEASE_20.1R3_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: 20.1R3-2021-5-10.0_RELEASE_20.1R3_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.

Junos OS Release Notes for vRR

IN THIS SECTION

These release notes accompany Junos OS Release 20.1R3 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

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

To learn about common BGP or routing Junos features supported on vRR for Junos OS 20.1R3, see What's New for MX Series routers.

What's New in Release 20.1R3

There are no new features for vRR in Junos OS Release 20.1R3.

What's New in Release 20.1R2

There are no new features for vRR in Junos OS Release 20.1R2.

What's New in Release 20.1R1

There are no new features for vRR in Junos OS Release 20.1R1.

What's Changed

IN THIS SECTION

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

To learn more about common BGP or routing changes in behavior or syntax in Junos OS 20.1R3, see What's Changed for MX Series routers.

What's Changed in Release 20.1R3

There are no changes in behavior or syntax for vRR in Junos OS Release 20.1R3.

What's Changed in Release 20.1R2

There are no changes in behavior or syntax for vRR in Junos OS Release 20.1R2.

What's Changed in Release 20.1R1

There are no changes in behavior or syntax for vRR in Junos OS Release 20.1R1.

Known Limitations

Learn about known limitations in this release for vRR.

To learn more about common BGP or routing known limitation in Junos OS 20.1R3, see Known Limitations 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.

There are no known behaviors or limitations for vRR in Junos OS Release 20.1R3.

Open Issues

Learn about open issues in this release for vRR.

To learn more about common BGP or routing open issues in Junos OS 20.1R3, see Open Issues 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.

There are no open issues for vRR in Junos OS Release 20.1R3.

Resolved Issues

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

To learn more about common BGP or routing resolved issues in Junos OS 20.1R3, see Resolved Issues 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: 20.1R3

Routing Protocols

- If output-queue-priority expedited update-tokens is configured, rpd might crash might upon BGP flapping. PR1545837
- Six PE device prefixes might not be removed from RIB upon reception of withdrawal from a BGP neighbor when the RIB sharding is enabled. PR1556271

Resolved Issues: 20.1R2

There are no resolved issues for vRR in Junos OS Release 20.1R2.

Resolved Issues: 20.1R1

There are no resolved issues for vRR in Junos OS Release 20.1R1.

Upgrading Using ISSU

In-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

For additional information about using ISSU on routing and switching devices, see the High Availability User Guide.

For additional information about using ISSU on security devices, see the Chassis Cluster User Guide for SRX Series Devices.

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

Licensing

Starting in 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 we've developed at Juniper Networks over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information about the list of supported products, see Juniper Flex Program.

Compliance Advisor

For regulatory compliance information about Common Criteria, FIPS, Homologation, RoHS2, and USGv6 for Juniper Networks products, see the Juniper Networks Compliance Advisor.

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you in exploring software feature information to find the right software release and product for your network. https://apps.juniper.net/feature-explorer/
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. prsearch.juniper.net.
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. apps.juniper.net/hct/home

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

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about Common Criteria, FIPS, Homologation, RoHS2, and USGv6 for Juniper Networks products. apps.juniper.net/compliance/.

Documentation Feedback

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

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

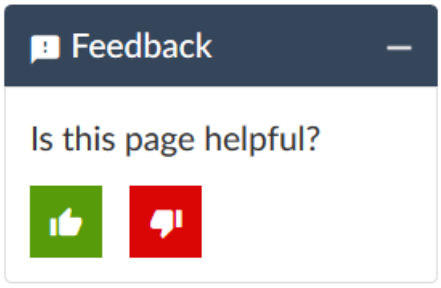

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active 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

12 May 2022—Revision 10, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

5 May 2022—Revision 9, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

24 March 2022—Revision 8, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

10 December 2021—Revision 7, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

7 October 2021—Revision 6, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

9 September 2021—Revision 5, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

5 August 2021—Revision 4, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

15 July 2021—Revision 3, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

8 July 2021—Revision 2, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

3 June 2021—Revision 1, Junos OS Release 20.1R3– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vSRX, and vRR.

22 April 2021—Revision 6, Junos OS Release 20.1R2– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX.

1 April 2021—Revision 5, Junos OS Release 20.1R2– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX.

11 February 2021—Revision 4, Junos OS Release 20.1R2– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX.

13 January 2021—Revision 3, Junos OS Release 20.1R2– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX.

3 December 2020—Revision 2, Junos OS Release 20.1R2– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX.

2 December 2020—Revision 1, Junos OS Release 20.1R2– ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX.

3 July 2020—Revision 7, Junos OS Release 20.1R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 June 2020—Revision 6, Junos OS Release 20.1R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

8 June 2020—Revision 5, Junos OS Release 20.1R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 May 2020—Revision 4, Junos OS Release 20.1R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

16 April 2020—Revision 3, Junos OS Release 20.1R1– PTX10003, PTX10008 Routers and the QFX5220 Switch.

3 April 2020—Revision 2, Junos OS Release 20.1R1– ACX Series, EX Series, MX Series, NFX Series,PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 March 2020—Revision 1, Junos OS Release 20.1R1– ACX Series, EX Series, MX Series, NFX Series,PTX Series, QFX Series, SRX Series, and Junos Fusion.