

Release Notes: Junos[®] OS Release 20.1R2 for the ACX Series, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX

22 April 2021

Contents [Introduction | 10](#)

[Junos OS Release Notes for ACX Series | 10](#)

[What's New | 11](#)

[What's New in Release 20.1R2 | 11](#)

[What's New in Release 20.1R1 | 11](#)

[What's Changed | 14](#)

[What's Changed in 20.1R2 | 14](#)

[What's Changed in 20.1R1 | 14](#)

[Known Limitations | 15](#)

[General Routing | 15](#)

[Open Issues | 16](#)

[General Routing | 16](#)

[Platform and Infrastructure | 18](#)

[Virtual Chassis | 18](#)

[Resolved Issues | 18](#)

[Resolved Issues: 20.1R2 | 19](#)

[Resolved Issues: 20.1R1 | 21](#)

[Documentation Updates | 23](#)

[Dynamic Host Configuration Protocol \(DHCP\) | 24](#)

Migration, Upgrade, and Downgrade Instructions | 24

Upgrade and Downgrade Support Policy for Junos OS Releases | 24

Junos OS Release Notes for EX Series Switches | 25

What's New | 26

What's New in Release 20.1R2 | 26

What's New in Release 20.1R1 | 27

What's Changed | 31

What's Changed in Release 20.1R2 | 31

What's Changed in Release 20.1R1 | 32

Known Limitations | 33

General Routing | 33

EVPN | 33

Infrastructure | 33

Platform and Infrastructure | 33

Open Issues | 34

General Routing | 34

High Availability (HA) and Resiliency | 35

Infrastructure | 35

Interfaces and Chassis | 36

Junos Fusion Provider Edge | 36

Layer 2 Ethernet Services | 36

Layer 2 Features | 36

Platform and Infrastructure | 36

Routing Protocols | 37

Resolved Issues | 37

Resolved Issues: 20.1R2 | 37

Resolved Issues: 20.1R1 | 42

Documentation Updates | 47

Dynamic Host Configuration Protocol (DHCP) | 47

Migration, Upgrade, and Downgrade Instructions | 47

Upgrade and Downgrade Support Policy for Junos OS Releases | 48

Junos OS Release Notes for JRR Series | 49

What's New | 49

What's Changed | 50

Known Limitations 50
Open Issues 50
Resolved Issues 51
Resolved Issues: 20.1R2 51
Resolved Issues: 20.1R1 52
Documentation Updates 52
Migration, Upgrade, and Downgrade Instructions 52
Upgrade and Downgrade Support Policy for Junos OS Releases 53
Junos OS Release Notes for Junos Fusion Enterprise 54
What's New 54
What's Changed 55
Known Limitations 55
Open Issues 56
Junos Fusion for Enterprise 56
Resolved Issues 57
Resolved Issues: 20.1R2 57
Resolved Issues: 20.1R1 57
Documentation Updates 58
Migration, Upgrade, and Downgrade Instructions 58
Basic Procedure for Upgrading Junos OS on an Aggregation Device 59
Upgrading an Aggregation Device with Redundant Routing Engines 61
Preparing the Switch for Satellite Device Conversion 61
Converting a Satellite Device to a Standalone Switch 62
Upgrade and Downgrade Support Policy for Junos OS Releases 63
Downgrading from Junos OS 63
Junos OS Release Notes for Junos Fusion Provider Edge 64
What's New 64
What's New in 20.1R2 Release 65
What's New in 20.1R1 Release 65
What's Changed 65
Known Limitations 66
Open Issues 66
Junos Fusion Provider Edge 67

Resolved Issues | 67

Resolved Issues: Release 20.1R2 | 68

Resolved Issues: Release 20.1R1 | 68

Documentation Updates | 68

Migration, Upgrade, and Downgrade Instructions | 69

Basic Procedure for Upgrading an Aggregation Device | 69

Upgrading an Aggregation Device with Redundant Routing Engines | 72

Preparing the Switch for Satellite Device Conversion | 72

Converting a Satellite Device to a Standalone Device | 74

Upgrading an Aggregation Device | 76

Upgrade and Downgrade Support Policy for Junos OS Releases | 76

Downgrading from Junos OS Release 20.1 | 77

Junos OS Release Notes for MX Series 5G Universal Routing Platform | 77

What's New | 78

What's New in 20.1R2 Release | 78

What's New in 20.1R1 Release | 78

What's Changed | 108

What's Changed in 20.1R2 | 109

What's Changed in 20.1R1 | 112

Known Limitations | 114

General Routing | 114

Infrastructure | 116

MPLS | 116

Platform and Infrastructure | 116

Services Applications | 117

Subscriber Management and Services | 117

Open Issues | 117

Class of Service (CoS) | 118

EVPN | 118

Forwarding and Sampling | 119

General Routing | 119

High Availability (HA) and Resiliency | 124

Infrastructure | 124

Interfaces and Chassis | 125

Layer 2 Ethernet Services 125
MPLS 126
Platform and Infrastructure 126
Routing Policy and Firewall Filters 128
Routing Protocols 128
User Interface and Configuration 129
VPNs 129
Resolved Issues 130
Resolved Issues: 20.1R2 130
Resolved Issues: 20.1R1 149
Documentation Updates 165
Dynamic Host Configuration Protocol (DHCP) 165
Migration, Upgrade, and Downgrade Instructions 166
Basic Procedure for Upgrading to Release 20.1R2 167
Procedure to Upgrade to FreeBSD 11.x based Junos OS 167
Procedure to Upgrade to FreeBSD 6.x based Junos OS 170
Upgrade and Downgrade Support Policy for Junos OS Releases 171
Upgrading a Router with Redundant Routing Engines 172
Downgrading from Release 20.1R2 172
Junos OS Release Notes for NFX Series 173
What's New 173
What's New in Release 20.1R2 174
What's New in Release 20.1R1 174
What's Changed 175
Known Limitations 176
Open Issues 176
Interfaces 177
Platform and Infrastructure 177
Virtual Network Functions (VNFs) 177
Resolved Issues 177
Resolved Issues: 20.1R2 178
Resolved Issues: 20.1R1 179
Documentation Updates 180

Migration, Upgrade, and Downgrade Instructions | 181

 Upgrade and Downgrade Support Policy for Junos OS Releases | 181

 Basic Procedure for Upgrading to Release 20.1 | 181

Junos OS Release Notes for PTX Series Packet Transport Routers | 183

 What's New | 184

 What's New in 20.1R2 | 184

 What's New in 20.1R1 | 184

 What's Changed | 191

 What's Changed in 20.1R2 Release | 191

 Known Limitations | 192

 General Routing | 193

 MPLS | 194

 Open Issues | 194

 General Routing | 195

 Infrastructure | 196

 MPLS | 196

 Routing Protocols | 196

 Resolved Issues | 197

 Resolved Issues: 20.1R2 | 197

 Resolved Issues: 20.1R1 | 199

 Documentation Updates | 201

 Dynamic Host Configuration Protocol (DHCP) | 201

Migration, Upgrade, and Downgrade Instructions | 202

 Basic Procedure for Upgrading to Release 20.1 | 202

 Upgrade and Downgrade Support Policy for Junos OS Releases | 205

 Upgrading a Router with Redundant Routing Engines | 205

Junos OS Release Notes for the QFX Series | 206

 What's New | 206

 What's New in Release 20.1R2 | 207

 What's New in Release 20.1R1 | 208

 What's Changed | 212

 What's Changed in 20.1R2 | 213

 What's Changed in 20.1R1 | 214

Known Limitations | 216

[Class of Service \(CoS\) | 217](#)

[EVPN | 217](#)

[Infrastructure | 217](#)

[Layer 2 Ethernet Services | 217](#)

[Layer 2 Features | 217](#)

[Platform and Infrastructure | 217](#)

Open Issues | 219

[Class of Service \(CoS\) | 220](#)

[EVPN | 220](#)

[High Availability \(HA\) and Resiliency | 220](#)

[Interfaces and Chassis | 220](#)

[Layer 2 Ethernet Services | 221](#)

[Layer 2 Features | 221](#)

[Platform and Infrastructure | 221](#)

[Routing Protocols | 223](#)

[Virtual Chassis | 224](#)

Resolved Issues | 224

[Resolved Issues: 20.1R2 | 224](#)

[Resolved Issues: 20.1R1 | 230](#)

Documentation Updates | 237

[Dynamic Host Configuration Protocol \(DHCP\) | 237](#)

Migration, Upgrade, and Downgrade Instructions | 238

[Upgrading Software on QFX Series Switches | 238](#)

[Installing the Software on QFX10002-60C Switches | 241](#)

[Installing the Software on QFX10002 Switches | 241](#)

[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 | 242](#)

[Installing the Software on QFX10008 and QFX10016 Switches | 244](#)

[Performing a Unified ISSU | 248](#)

[Preparing the Switch for Software Installation | 249](#)

[Upgrading the Software Using Unified ISSU | 249](#)

[Upgrade and Downgrade Support Policy for Junos OS Releases | 251](#)

[Junos OS Release Notes for SRX Series | 252](#)

[What's New | 253](#)

[Release 20.1R2 New and Changed Features | 253](#)

[Release 20.1R1 New and Changed Features | 253](#)

[What's Changed | 260](#)

[What's Changed in Release 20.1R2 | 261](#)

[What's Changed in Release 20.1R1 | 262](#)

[Known Limitations | 264](#)

[J-Web | 265](#)

[Platform and Infrastructure | 265](#)

[VPNs | 265](#)

[Open Issues | 266](#)

[ATP Cloud | 266](#)

[Flow-Based and Packet-Based Processing | 266](#)

[J-Web | 267](#)

[Routing Policy and Firewall Filters | 267](#)

[VPNs | 267](#)

[Resolved Issues | 268](#)

[Resolved Issues: 20.1R2 | 268](#)

[Resolved Issues: 20.1R1 | 273](#)

[Documentation Updates | 278](#)

[Dynamic Host Configuration Protocol \(DHCP\) | 279](#)

[Migration, Upgrade, and Downgrade Instructions | 279](#)

[Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 279](#)

[Junos OS Release Notes for vMX | 280](#)

[What's New | 281](#)

[What's Changed | 281](#)

[Known Limitations | 281](#)

[Open Issues | 281](#)

[Resolved Issues | 281](#)

[Platform and Infrastructure | 282](#)

[Licensing | 282](#)

[Upgrade Instructions | 283](#)

[Junos OS Release Notes for vSRX | 283](#)

[What's New | 283](#)

[What's New in Release 20.1R2 | 284](#)

[What's Changed | 284](#)

[What's Changed in Release 20.1R2 | 285](#)

[Known Limitations | 285](#)

[Flow-Based and Packet-Based Processing | 286](#)

[J-Web | 286](#)

[User Access and Authentication | 286](#)

[Open Issues | 286](#)

[Intrusion Detection and Prevention \(IDP\) | 287](#)

[Resolved Issues | 287](#)

[Resolved Issues: 20.1R2 | 287](#)

[Migration, Upgrade, and Downgrade Instructions | 289](#)

[Upgrading Software Packages | 290](#)

[Validating the OVA Image | 295](#)

[Upgrading Using ISSU | 296](#)

[Licensing | 296](#)

[Compliance Advisor | 297](#)

[Finding More Information | 297](#)

[Documentation Feedback | 297](#)

[Requesting Technical Support | 299](#)

[Self-Help Online Tools and Resources | 299](#)

[Creating a Service Request with JTAC | 300](#)

[Revision History | 300](#)

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, and vSRX.

These release notes accompany Junos OS Release 20.1R2 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, and vSRX. 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 296](#)
 - [Licensing on page 296](#)
 - [Compliance Advisor on page 297](#)
 - [Finding More Information on page 297](#)
 - [Documentation Feedback on page 297](#)
 - [Requesting Technical Support on page 299](#)

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [What's New | 11](#)
- [What's Changed | 14](#)
- [Known Limitations | 15](#)
- [Open Issues | 16](#)
- [Resolved Issues | 18](#)

- Documentation Updates | [23](#)
- Migration, Upgrade, and Downgrade Instructions | [24](#)

These release notes accompany Junos OS Release 20.1R2 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

- [What's New in Release 20.1R2 | 11](#)
- [What's New in Release 20.1R1 | 11](#)

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

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 device 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 | 14](#)

[Known Limitations | 15](#)

[Open Issues | 16](#)

[Resolved Issues | 18](#)

[Documentation Updates | 23](#)

[Migration, Upgrade, and Downgrade Instructions | 24](#)

What's Changed

IN THIS SECTION

- [What's Changed in 20.1R2 | 14](#)
- [What's Changed in 20.1R1 | 14](#)

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

What's Changed in 20.1R2

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](#) .]

What's Changed in 20.1R1

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

SEE ALSO

[What's New | 11](#)

[Known Limitations | 15](#)

[Open Issues | 16](#)

[Resolved Issues | 18](#)

[Documentation Updates | 23](#)

[Migration, Upgrade, and Downgrade Instructions | 24](#)

Known Limitations

IN THIS SECTION

- General Routing | 15

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](#)

SEE ALSO

-
- [What's New | 11](#)
 - [What's Changed | 14](#)
 - [Open Issues | 16](#)
 - [Resolved Issues | 18](#)
 - [Documentation Updates | 23](#)
 - [Migration, Upgrade, and Downgrade Instructions | 24](#)
-

Open Issues

IN THIS SECTION

- [General Routing | 16](#)
- [Platform and Infrastructure | 18](#)
- [Virtual Chassis | 18](#)

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

- Several error logs are observed when the link in the LSP primary path flaps. [PR1204714](#)
- Loopback status is not displayed for ot interfaces on the CLI, which is available from vty only. [PR1358017](#)
- The SD threshold is set above the SF threshold. [PR1376869](#)
- On the ACX6360-OR router, enhancement is needed for the FRR BER threshold SNMP support. [PR1383303](#)
- On the ACX5448 router with 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet interfaces, link fault signaling does not work. [PR1401718](#)
- A **jnxIfOtnOperState** trap notification is sent for all ot- interfaces. [PR1406758](#)
- The PTP client gets stuck in the **Initialize** state intermittently due to the ARP resolution issue when PTP and interfaces flap continuously. [PR1410746](#)

- Traffic loss and SNMP slow response issues are seen if an optical transceiver is removed and then inserted back into the same interface. [PR1418696](#)
- On the ACX5000 router, high CPU usage on fxpc process might be observed. [PR1419761](#)
- The FPC might crash with the em2 configuration during initialization. [PR1429212](#)
- Protocols get forwarded when you use the nonexisting SSM map source address in IGMPv3 instead of pruning. [PR1435648](#)
- Memory leak is observed while running the longevity check. [PR1438358](#)
- RED drops might be seen after link flaps or CoS configuration changes. [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 IPv6 BFD sessions flap when configured below 100 ms. [PR1456237](#)
- MPLS packet load balancing is performed without enabling hashing. [PR1475363](#)
- The FPC process might continuously crash after deactivating or activating the loopback filter or rebooting 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 the PTP environment, some vendor devices acting as client expect announced messages at an interval of -3 (8 pps) from the upstream primary device. [PR1507782](#)
- ACX1100, ACX2100, ACX2200, ACX2000, and ACX4000 routers might stop forwarding transit and control traffic. [PR1508534](#)
- 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](#)
- On the ACX5448 routers, the OSPF state is not as expected. [PR1543667](#)
- The ACX5448 router as transit for BGP labeled unicast drops traffic. [PR1547713](#)
- 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](#)
- On the ACX5448 routers, the transit DHCP packets drop is observed. [PR1517420](#)
- On the ACX500-I routers, the **show services session count** command does not work as expected. [PR1520305](#)

Platform and Infrastructure

- The CFM REMOTE MEP does not come up after configuration or if the MEP remains in the **Start** state. [PR1460555](#)

Virtual Chassis

- In the ACX5000 router, the following false positive parity error messages are observed:
`_soc_mem_array_sbisdma_read`. [PR1276970](#)

SEE ALSO

[What's New | 11](#)

[What's Changed | 14](#)

[Known Limitations | 15](#)

[Resolved Issues | 18](#)

[Documentation Updates | 23](#)

[Migration, Upgrade, and Downgrade Instructions | 24](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.1R2 | 19](#)
- [Resolved Issues: 20.1R1 | 21](#)

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: 20.1R2

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-signalized 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 traffic to 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 asynchronized 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

[What's New | 11](#)

[What's Changed | 14](#)

[Known Limitations | 15](#)

[Open Issues | 16](#)

[Documentation Updates | 23](#)

[Migration, Upgrade, and Downgrade Instructions | 24](#)

Documentation Updates

IN THIS SECTION

- [Dynamic Host Configuration Protocol \(DHCP\) | 24](#)

This section lists the errata and changes in Junos OS Release 20.1R2 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

[What's New | 11](#)

[What's Changed | 14](#)

[Known Limitations | 15](#)

[Open Issues | 16](#)

[Resolved Issues | 18](#)

[Migration, Upgrade, and Downgrade Instructions | 24](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 24](#)

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 | 11](#)
- [What's Changed | 14](#)
- [Known Limitations | 15](#)
- [Open Issues | 16](#)
- [Resolved Issues | 18](#)
- [Migration, Upgrade, and Downgrade Instructions | 24](#)

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- [What's New | 26](#)
- [What's Changed | 31](#)
- [Known Limitations | 33](#)
- [Open Issues | 34](#)
- [Resolved Issues | 37](#)
- [Documentation Updates | 47](#)
- [Migration, Upgrade, and Downgrade Instructions | 47](#)

These release notes accompany Junos OS Release 20.1R2 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

- [What's New in Release 20.1R2 | 26](#)
- [What's New in Release 20.1R1 | 27](#)

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 20.1R2: EX2300, EX2300-C, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

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 | 31](#)

[Known Limitations | 33](#)

[Open Issues | 34](#)

[Resolved Issues | 37](#)

[Documentation Updates | 47](#)

[Migration, Upgrade, and Downgrade Instructions | 47](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.1R2 | 31](#)
- [What's Changed in Release 20.1R1 | 32](#)

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

What's Changed in Release 20.1R2

There are no changes in behavior and syntax for EX Series in Junos OS Release 20.1R2.

What's Changed in Release 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

- **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](#).]

SEE ALSO

[What's New | 26](#)

[Known Limitations | 33](#)

[Open Issues | 34](#)

[Resolved Issues | 37](#)

[Documentation Updates | 47](#)

[Migration, Upgrade, and Downgrade Instructions | 47](#)

Known Limitations

IN THIS SECTION

- General Routing | [33](#)
- EVPN | [33](#)
- Infrastructure | [33](#)
- Platform and Infrastructure | [33](#)

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](#)

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](#)

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 | 26](#)

[What's Changed | 31](#)

[Open Issues | 34](#)

[Resolved Issues | 37](#)

[Documentation Updates | 47](#)

[Migration, Upgrade, and Downgrade Instructions | 47](#)

Open Issues

IN THIS SECTION

- [General Routing | 34](#)
- [High Availability \(HA\) and Resiliency | 35](#)
- [Infrastructure | 35](#)
- [Interfaces and Chassis | 36](#)
- [Junos Fusion Provider Edge | 36](#)
- [Layer 2 Ethernet Services | 36](#)
- [Layer 2 Features | 36](#)
- [Platform and Infrastructure | 36](#)
- [Routing Protocols | 37](#)

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.

General Routing

- On the EX9208 device, 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 the Strict mode does not work. [PR1417546](#)
- Legacy EX Series platforms might reboot when removing or adding `input egress all` for the analyzer.
[PR1448123](#)
- On the EX9214 device, the following error message is observed after reboot and MACsec-enabled link flaps: `errorlib_set_error_log(): err_id(-1718026239).` [PR1448368](#)
- On the EX4300 Virtual Chassis, a switch might drop the Dot1x client TLS packet. [PR1464365](#)

- The following error message might appear: **Failed to complete DFE tuning**. This error message has no functional impact and can be ignored. [PR1473280](#)
- Classifiers binding applied on a wildcard gets overwritten by a different classifier type when applied on a single interface. [PR1490699](#)
- While verifying the **Last-change op-state** value through XML, the **rpc-reply** message is inappropriate. [PR1492449](#)
- 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](#)
- On the EX3400 and EX2300 switches, the upgrade fails due to the lack of available storage. [PR1539293](#)
- Domain and VLAN parameters are not as expected while verifying the **show dot1x interface detail** command. [PR1553596](#)
- OSPF and OSPF3 adjacencies uptime is more than expected after NSSU upgrade and outage is higher than the expected. [PR1551925](#)
- 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](#)
- The DHCP traffic might not be forwarded correctly while sending the DHCP unicast packets. [PR1512175](#)
- On the EX4300 device, script fails while committing the IPSEC authentication configuration due to the missing **algorithm** statement. [PR1557216](#)

High Availability (HA) and Resiliency

- On the EX9204 devices, the `vmcore.live` process generates core file with traces `0xffffffff80590e97` in `timercb`, `0xffffffff805ba5c1` in `lapi_handle_timer`, `0xffffffff80543df3` in `uma_zfree_arg` on the AD or SD setup. [PR1557571](#)

Infrastructure

- Some PIM groups are not able to send out native multicast traffic. [PR1209585](#)
- HSRPv2 IPv6 packets might be dropped if IGMP snooping is enabled. [PR1232403](#)
- The FPC crashes with pfem generating a core file if a large-scale number of firewall filters are configured. [PR1434927](#)
- The following error message is observed continuously in the AD with base configurations: **IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151) failed.** [PR1485038](#)
- On an EX4300-MP device, 9000 IPv6 MC routes can be installed. If more IPv6 MC routes are added, error messages are displayed. [PR1493671](#)

Interfaces and Chassis

- The GRES and VSTP ports cost on the aggregated Ethernet interfaces might get changed, leading to a topology change. [PR1174213](#)

Junos Fusion Provider Edge

- On a 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 because a few packets transmitted from the egress of AD1 are 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 packets are 122 bytes. [PR1450373](#)

Layer 2 Ethernet Services

- The DHCP decline packets are not forwarded to the DHCP server when forward-only is set within dhcp-reply. [PR1429456](#)

Layer 2 Features

- G-ARPs are sent from the switch once in every 10 minutes. [PR1192520](#)

Platform and Infrastructure

- On the EX9208 devices, traffic loss is observed if ingress and egress ports are in different FPCs. [PR1429714](#)
- On the EX9208 device, 33 percent degradation with MAC learning rate is observed in Junos OS Release 19.3R1 compared to Junos OS Release 18.4R1. [PR1450729](#)
- The pfex_junos process generates core file at 0x01847994 in `pfeman_watchdog (arg=< optimized out>)` at `../../../../src/pfe/common/applications/pfeman/pfeman_rt_pfex.c:1411`. [PR1535178](#)
- The BGP session replication might fail to start after the session crashes on a backup Routing Engine. [PR1552603](#)

Routing Protocols

- Sending multicast traffic to the downstream receiver on the Trio based Virtual Chassis might fail. [PR1555518](#)

SEE ALSO

[What's New | 26](#)

[What's New | 26](#)

[Known Limitations | 33](#)

[Resolved Issues | 37](#)

[Documentation Updates | 47](#)

[Migration, Upgrade, and Downgrade Instructions | 47](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.1R2 | 37](#)
- [Resolved Issues: 20.1R1 | 42](#)

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: 20.1R2

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 statistics** 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-It 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 dhcpcd 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 | 26](#)

[What's Changed | 31](#)

[Known Limitations | 33](#)

[Open Issues | 34](#)

[Documentation Updates | 47](#)

[Migration, Upgrade, and Downgrade Instructions | 47](#)

Documentation Updates

IN THIS SECTION

- [Dynamic Host Configuration Protocol \(DHCP\) | 47](#)

This section lists the errata and changes in Junos OS Release 20.1R2 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 | 37](#)
 - [What's Changed | 31](#)
 - [Known Limitations | 33](#)
 - [Open Issues | 34](#)
 - [Resolved Issues | 37](#)
 - [Migration, Upgrade, and Downgrade Instructions | 47](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 48](#)

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 | 26](#)

[What's Changed | 31](#)

[Known Limitations | 33](#)

[Open Issues | 34](#)

[Resolved Issues | 37](#)

[Documentation Updates | 47](#)

Junos OS Release Notes for JRR Series

IN THIS SECTION

- [What's New | 49](#)
- [What's Changed | 50](#)
- [Known Limitations | 50](#)
- [Open Issues | 50](#)
- [Resolved Issues | 51](#)
- [Documentation Updates | 52](#)
- [Migration, Upgrade, and Downgrade Instructions | 52](#)

These release notes accompany Junos OS Release 20.1R2 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 20.1R2.

SEE ALSO

- [What's Changed | 50](#)
- [Known Limitations | 50](#)
- [Open Issues | 50](#)
- [Resolved Issues | 51](#)
- [Documentation Updates | 52](#)
- [Migration, Upgrade, and Downgrade Instructions | 52](#)

What's Changed

There are no changes in behavior and syntax for JRR Series in Junos OS Release 20.1R2.

SEE ALSO

- [What's New | 49](#)
- [Known Limitations | 50](#)
- [Open Issues | 50](#)
- [Resolved Issues | 51](#)
- [Documentation Updates | 52](#)
- [Migration, Upgrade, and Downgrade Instructions | 52](#)

Known Limitations

There are no known limitations for JRR Series in Junos OS Release 20.1R2.

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 | 49](#)
- [What's Changed | 50](#)
- [Open Issues | 50](#)
- [Resolved Issues | 51](#)
- [Documentation Updates | 52](#)
- [Migration, Upgrade, and Downgrade Instructions | 52](#)

Open Issues

There are no open issues for JRR Series in Junos OS 20.1R1 Release.

SEE ALSO

-
- [What's New | 49](#)
 - [What's Changed | 50](#)
 - [Known Limitations | 50](#)
 - [Resolved Issues | 51](#)
 - [Documentation Updates | 52](#)
 - [Migration, Upgrade, and Downgrade Instructions | 52](#)
-

Resolved Issues

IN THIS SECTION

-  [Resolved Issues: 20.1R2 | 51](#)
-  [Resolved Issues: 20.1R1 | 52](#)

Learn about resolved issues for JRR Series in Junos OS 20.1R2 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: 20.1R2

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 | 49](#)
- [What's Changed | 50](#)
- [Known Limitations | 50](#)
- [Open Issues | 50](#)
- [Documentation Updates | 52](#)
- [Migration, Upgrade, and Downgrade Instructions | 52](#)

Documentation Updates

There are no errata or changes in Junos OS Release 20.1R2 documentation for JRR200 Route Reflectors.

SEE ALSO

- [What's New | 49](#)
- [What's Changed | 50](#)
- [Known Limitations | 50](#)
- [Open Issues | 50](#)
- [Resolved Issues | 51](#)
- [Migration, Upgrade, and Downgrade Instructions | 52](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 53](#)

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 | 49](#)

[What's Changed | 50](#)

[Known Limitations | 50](#)

[Resolved Issues | 51](#)

[Open Issues | 50](#)

[Documentation Updates | 52](#)

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

- [What's New | 54](#)
- [What's Changed | 55](#)
- [Known Limitations | 55](#)
- [Open Issues | 56](#)
- [Resolved Issues | 57](#)
- [Documentation Updates | 58](#)
- [Migration, Upgrade, and Downgrade Instructions | 58](#)

These release notes accompany Junos OS Release 20.1R2 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 20.1R2 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 | 55](#)
 - [Known Limitations | 55](#)
 - [Open Issues | 56](#)
 - [Resolved Issues | 57](#)
 - [Documentation Updates | 58](#)
 - [Migration, Upgrade, and Downgrade Instructions | 58](#)

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 20.1R2 for Junos fusion for enterprise.

SEE ALSO

-
- [What's New | 54](#)
 - [Known Limitations | 55](#)
 - [Open Issues | 56](#)
 - [Resolved Issues | 57](#)
 - [Documentation Updates | 58](#)
 - [Migration, Upgrade, and Downgrade Instructions | 58](#)

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.1R2 for Junos fusion for enterprise.

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

SEE ALSO

-
- [What's New | 54](#)

[What's Changed | 55](#)

[Open Issues | 56](#)

[Resolved Issues | 57](#)

[Documentation Updates | 58](#)

[Migration, Upgrade, and Downgrade Instructions | 58](#)

Open Issues

IN THIS SECTION

- [Junos Fusion for Enterprise | 56](#)

This section lists the known issues in hardware and software in Junos OS Release 20.1R2 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.

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 | 54](#)

[What's Changed | 55](#)

[Known Limitations | 55](#)

[Resolved Issues | 57](#)

[Documentation Updates | 58](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.1R2 | 57](#)
- [Resolved Issues: 20.1R1 | 57](#)

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: 20.1R2

- 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 Changed | 55](#)

[Known Limitations | 55](#)

[Open Issues | 56](#)

[What's Changed | 55](#)

[Migration, Upgrade, and Downgrade Instructions | 58](#)

Documentation Updates

There are no errata or changes in Junos OS Release 20.1R2 for documentation for Junos fusion for enterprise.

SEE ALSO

[What's New | 54](#)

[What's Changed | 55](#)

[Known Limitations | 55](#)

[Open Issues | 56](#)

[Resolved Issues | 57](#)

[Migration, Upgrade, and Downgrade Instructions | 58](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device | 59](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 61](#)
- [Preparing the Switch for Satellite Device Conversion | 61](#)
- [Converting a Satellite Device to a Standalone Switch | 62](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 63](#)
- [Downgrading from Junos OS | 63](#)

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 | 54](#)

[What's Changed | 55](#)

[Known Limitations | 55](#)

[Open Issues | 56](#)

[Resolved Issues | 57](#)

[Documentation Updates | 58](#)

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- [What's New | 64](#)
- [What's Changed | 65](#)
- [Known Limitations | 66](#)
- [Open Issues | 66](#)
- [Resolved Issues | 67](#)
- [Documentation Updates | 68](#)
- [Migration, Upgrade, and Downgrade Instructions | 69](#)

These release notes accompany Junos OS Release 20.1R2 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

- [What's New in 20.1R2 Release | 65](#)
- [What's New in 20.1R1 Release | 65](#)

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

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 | 65](#)
 - [Known Limitations | 66](#)
 - [Open Issues | 66](#)
 - [Resolved Issues | 67](#)
 - [Documentation Updates | 68](#)
 - [Migration, Upgrade, and Downgrade Instructions | 69](#)

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

-
- [What's New | 64](#)
 - [Known Limitations | 66](#)
 - [Open Issues | 66](#)
 - [Resolved Issues | 67](#)
 - [Documentation Updates | 68](#)
 - [Migration, Upgrade, and Downgrade Instructions | 69](#)

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.1R2 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

[What's New | 64](#)

[What's Changed | 65](#)

[Open Issues | 66](#)

[Resolved Issues | 67](#)

[Documentation Updates | 68](#)

[Migration, Upgrade, and Downgrade Instructions | 69](#)

Open Issues

IN THIS SECTION

- [Junos Fusion Provider Edge | 67](#)

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 | 64](#)

[What's Changed | 65](#)

[Known Limitations | 66](#)

[Resolved Issues | 67](#)

[Documentation Updates | 68](#)

[Migration, Upgrade, and Downgrade Instructions | 69](#)

Resolved Issues

IN THIS SECTION

-  [Resolved Issues: Release 20.1R2 | 68](#)
-  [Resolved Issues: Release 20.1R1 | 68](#)

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 20.1R2

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 | 64](#)

[What's Changed | 65](#)

[Known Limitations | 66](#)

[Open Issues | 66](#)

[Documentation Updates | 68](#)

[Migration, Upgrade, and Downgrade Instructions | 69](#)

Documentation Updates

There are no errata or changes in Junos OS Release 20.1R2 documentation for Junos fusion for provider edge.

SEE ALSO

[What's New | 64](#)

[What's Changed | 65](#)

[Known Limitations | 66](#)

[Open Issues | 66](#)

[Resolved Issues | 67](#)

[Migration, Upgrade, and Downgrade Instructions | 69](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | [69](#)
- Upgrading an Aggregation Device with Redundant Routing Engines | [72](#)
- Preparing the Switch for Satellite Device Conversion | [72](#)
- Converting a Satellite Device to a Standalone Device | [74](#)
- Upgrading an Aggregation Device | [76](#)
- Upgrade and Downgrade Support Policy for Junos OS Releases | [76](#)
- Downgrading from Junos OS Release 20.1 | [77](#)

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 20.1R2 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/jinstall64-20.1R2.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot  
source/jinstall-20.1R2.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/jinstall64-20.1R2.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot  
source/jinstall-20.1R2.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 20.1R2 **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.

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, unplug 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 20.1R2, 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 | 64](#)

[What's Changed | 65](#)

[Known Limitations | 66](#)

[Open Issues | 66](#)

[Resolved Issues | 67](#)

[Documentation Updates | 68](#)

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

IN THIS SECTION

- [What's New | 78](#)
- [What's Changed | 108](#)
- [Known Limitations | 114](#)
- [Open Issues | 117](#)
- [Resolved Issues | 130](#)
- [Documentation Updates | 165](#)
- [Migration, Upgrade, and Downgrade Instructions | 166](#)

These release notes accompany Junos OS Release 20.1R2 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

- [What's New in 20.1R2 Release | 78](#)
- [What's New in 20.1R1 Release | 78](#)

Learn about new features introduced in Junos OS Release 20.1R2 for MX Series routers.

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 reach 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 default 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 device 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 support selective 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-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 | 108](#)

[Known Limitations | 114](#)

[Open Issues | 117](#)

[Resolved Issues | 130](#)

[Documentation Updates | 165](#)

[Migration, Upgrade, and Downgrade Instructions | 166](#)

What's Changed

IN THIS SECTION

- [What's Changed in 20.1R2 | 109](#)
- [What's Changed in 20.1R1 | 112](#)

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

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 20.1R1

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

[What's New | 78](#)

[Known Limitations | 114](#)

[Open Issues | 117](#)

[Resolved Issues | 130](#)

[Documentation Updates | 165](#)

[Migration, Upgrade, and Downgrade Instructions | 166](#)

Known Limitations

IN THIS SECTION

- General Routing | [114](#)
- Infrastructure | [116](#)
- MPLS | [116](#)
- Platform and Infrastructure | [116](#)
- Services Applications | [117](#)
- Subscriber Management and Services | [117](#)

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](#)
- Dynamic SR-TE tunnels does not get automatically re-created at new master after Routing Engine switchover. [PR1474397](#)
- 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](#)
- 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](#)

MPLS

- 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](#)

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](#)

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 | 78](#)
- [What's Changed | 108](#)
- [Open Issues | 117](#)
- [Resolved Issues | 130](#)
- [Documentation Updates | 165](#)
- [Migration, Upgrade, and Downgrade Instructions | 166](#)

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 118](#)
- [EVPN | 118](#)
- [Forwarding and Sampling | 119](#)
- [General Routing | 119](#)

●	High Availability (HA) and Resiliency 124
●	Infrastructure 124
●	Interfaces and Chassis 125
●	Layer 2 Ethernet Services 125
●	MPLS 126
●	Platform and Infrastructure 126
●	Routing Policy and Firewall Filters 128
●	Routing Protocols 128
●	User Interface and Configuration 129
●	VPNs 129

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.

Class of Service (CoS)

- CoS EXP classifier and rewrite with the **mpls-inet-both-non-vpn** protocol option is not working as expected. [PR1479575](#)

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](#)
- The **no-arp-suppression** statement is required for MAC learning to happen across the EVPN domain on a static VTEP. [PR1517591](#)
- 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**. As a result, the data traffic sent via these EVPN routes does not push VLAN ID in the inner Ethernet header. This might result in traffic getting discarded on the remote PE device. [PR1547275](#)

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. [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](#)
- On MX platforms with EVPN-VXLAN implemented, ARP requests received on the VXLAN tunnel endpoint (VTEP) might not forward to CE device or proxy ARP role on VTEP might not work properly. This issue could happen when **no-arp-suppression** statement is disabled under EVPN instance combined with static VXLAN implementation. [PR1546631](#)

General Routing

- The fxp0 marked as **Dest-route-down** because of specific operations such as disabling and enabling operations. [PR1052725](#)
- On a vMX platform, the performance of the X710 NIC is lower compared to the performance of the 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, under the condition where corruption is with the host root file system, the node is booting with the previous vmhost software instead of booting from the alternate disk. [PR1281554](#)
- 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 Queue Telemetry Transport (MQTT) scenario, about 4000 KB of memory leakage might be seen every 30 seconds. However, on very long runs, this leakage uses up high memory, which can indirectly impact other running daemons. [PR1324531](#)
- With regard to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infrastructure relies on the integrity of the TCP connections and hence the reactions to failure situations might not be handled in a graceful way. For example, TCP connection timeout because of jlock hog crossing the

boundary value (5 seconds) can cause bad consequences for the MX Series Virtual Chassis. Currently, there is no other easy solution to reduce this jlock hog besides enabling the marker infrastructure in an MX Series Virtual Chassis setup. [PR1332765](#)

- On MX2010 and MX2020 routers equipped with SFB2, some error messages are seen in the logs. There is no operational impact caused by these messages. [PR1363587](#)
- A 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](#)
- The virtio throughput remains the same for multi-queue and single-queue deployments. [PR1389338](#)
- Traffic destined to the VRRP VIP gets dropped as the filter is not updated to the related logical interface. [PR1390367](#)
- FPC core files are generated on multiple additions or deletions of hierarchical CoS from pseudowire devices. As a workaround, remove the pseudowire device without changing the hierarchical CoS configuration. [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](#)
- With LACP enabled, deactivating a remote aggregate Ethernet member link makes the local member link move to LACP detached state and causes traffic drop on that member link. The same scenario happens when a new member link is added where the other end of that link is not yet configured with LACP. [PR1423707](#)
- 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](#)
- On MX Series platforms, if the clock frequency is slowly changing on CB0 (slow drift), the clock source for MPC-3D-16XGE-SFPP might 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. Even though the configuration gets committed, the feature does not work. [PR1435855](#)
- Interface hold-down timers cannot be achieved for less than 15 seconds on the MPC11E line card. [PR1444516](#)

- 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](#)
- On the MPC11E line card, the FIB download rates are lower than on the MPC10E by 30 percent. [PR1456816](#)
- 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 these 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](#)
- 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. [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](#)
- The following syslog error messages are harmless and expected during ISSU or GRES or FPC offline/online scenarios: [Oct 3 08:48:35.836 LOG: Err] ifl ps240.1 (1712): child ifl lt-1/0/0.32767 (7709) already there [Oct 3 08:48:35.836 LOG: Err] IFRT: 'Aggregate interface ifl add req' (opcode 87) failed [Oct 3 08:48:35.836 LOG: Err] ifl 1712, child ifl 7709; agg add failed. [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](#)
- Unable to get the service sessions when NAT64 is configured with the destination-prefix length 32. [PR1468058](#)
- The FPC might take additional time to come online during movement of MPC11 from one GNF to another GNF. [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](#)
- 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](#)
- Dynamic SR-TE tunnels do not get automatically re-created at the new primary Routing Engine after the Routing Engine switchover. [PR1474397](#)
- 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](#)
- 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](#)
- The BGP sessions over the PS interfaces anchored over RLT might flap during ISSU. [PR1478693](#)
- 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](#)
- Error message **fpc0 user.crit aftd-trio: [Critical] Em: Possible out of order deletion of AftNode #012#012#012 AftNode details - AftIndirect token:230791 group:0 nodeMask:0xfffffffffffffff indirect:333988 hwInstall:1#012** is seen while executing the IP/MPLS script in 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](#)
- 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 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](#)
- On an MX2020 device, the MPC11 line card is not supported. [PR1503605](#)
- 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](#)
- Disruptive switchover (no GRES or NSR configured) can lead to stale Periodic Packet Management (PPM) entries programmed on the new primary 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](#)
- **Show interfaces redundancy rlt0** shows current status as primary down as the FPC is still in ready state after the RLT failover. Issue occurs only with specific RLT interface configuration followed by FPC restart. In this PR, GRES was done when the FPC was doing the interface cleanup after FPC restart. The amount of time taken for interface cleanup on an FPC depends on scale. If GRES is done before FPC comes to clean state, the backup Routing Engine might not be in sync with the primary Routing Engine. This might create inconsistent interface states after GRES. [PR1518543](#)
- Inconsistent core.python2.7.mpc0 core files are seen. [PR1534568](#)
- If you pull out the MPC before the router is offline and the chassisd process is not able to process this event on the master Routing Engine because of additional mastership switch, and later the MPC which is pulled out is inserted back to the slot, many Switch Fabric Board (SFB) might go offline because of **max_total_cell_usage** overflow condition on the xfchip. MX2020 platform with SFB2 is not exposed to such event if the **disable-grant-bypass** is configured. [PR1535787](#)
- On MX platforms with Junos OS Release 18.1 or higher release, chassisd memory leak might be caused by configuration commit. When chassisd consumes ~3.4 GB of memory it might crash, chassisd crash might cause GRES or/and FPC to restart. If GRES is enabled, commits are being synchronized between Routing Engines, so backup Routing Engine chassisd might suffer from memory leak too. [PR1537194](#)
- If the Packet Forwarding Engine processes distributed IGMP pseudo logical interfaces is deleted, it attempts to delete all the associated multicast flows. On a scaled setup, deleting several thousand multicast flows hogs CPU for long time that it is killed by the scheduler, resulting in core file generation. This is a rare condition, seen only on scaled distributed IGMP setup. [PR1537846](#)
- On MX2020 routers with MPC11 line card, plane offline event is getting stuck. [PR1546449](#)

- In syslog output, sylog-local-tag name is truncated (it shows as SYSLOG_SF) when you configure the syslog local-log-tag as SYSLOG_SFW. [PR1547505](#)
- The following leak is observed during the period of churn for the sensor group bound to RSVP P2MP tunnels: **SENSOR APP DWORD**.[PR1547698](#)
- Validation of OCSP certificate might not go through for some CA servers. [PR1548268](#)
- Whenever AMS undergoes any update, flowd process crashes. When AMS sends LB blob, USF-HCM-PKT plugin consumes and frees it. But it frees in legacy platform (Non-USF) manner which leads flowd process to crash. Basically, USF-HCM-PKT plugin is not supported in the USF platform and it is supported only in the Non-USF platform alone. [PR1548928](#)
- In the ptpt freerun state, output of the **show ptpt clock** might continue to show the clock attributes of previously locked clock. [PR1553369](#)
- On MX240, MX480, and MX960 platforms, with default increased-bandwidth fabric mode and SCBE3, if we have MPC3 or MPC3-NG line card exist on the system along with high bandwidth MPC, during high traffic situation or bursty traffic through the fabric towards MPC3/MPC3-NG line card. MX fabric might report unreachable destination condition and causes fabric healing to trigger in. This issue is exacerbated when having MPC7 or MPC10 line cards installed due to high fabric bandwidth that can be generated. [PR1553641](#)

High Availability (HA) and Resiliency

- When an ZPL is done while traffic is running, some BGP sessions might flap, leading to traffic loss. The drop is transient and traffic recovers after the ZPL. [PR1487144](#)

Infrastructure

- The following error message might be seen after an upgrade: **invalid SMART checksum**. [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](#)
- Add 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](#)

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 user a 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](#)
- After Routing Engine switchover without nonstop routing (NSR) on the broadband network gateway (BNG), some VRF routing instances might experience silent dropping of traffic destined to the hosts behind a static PPPoE subscriber's CPE device. The affected routing instances are configured without the **vrf-table-label** statement and should have a static route configured with the pp0.xxx interface as a next hop. [PR1488302](#)
- On MX platforms, under subscriber environment and unnumbered IP address borrowed from the loopback interface configured on Demux interface and aggregated Ethernet interface, subscriber session might flap if the IP address of the loopback interface IP is changed. [PR1544257](#)
- After an ASIC error, the logical interface down messages were being sent to VC-MM. They should have been sent to the local chassis master where the error was reported. [PR1552588](#)
- The device control daemon (dcd) memory leak issue might be observed on pushing the scaled routing-instance configuration with bridge-domain stanza into the ephemeral database. [PR1553148](#)

Layer 2 Ethernet Services

- When you revert from an Enhanced Switch Control Board (SCBE) upgrade, the SCB fails with the following error message: CHASSISD_FASIC_PIO_READ_ERROR. [PR980340](#)
- The DHCP DECLINE packets are not forwarded to the DHCP server when **forward-only** is set within **dhcp-reply**. [PR1429456](#)

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](#)
- On all Junos OS platforms with RSVP-TE configured, when a transit router carries a large number of LSP's (for example, 60000 and higher) and all those LSPs undergo FRR (for example, when link carrying large number of LSP's flap) then, the rpd scheduler might slip and LDP session tear down might be observed after the link flap. [PR1516657](#)

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 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](#)
- 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](#)
- The CFM remote MEP is not coming up after configuration or remains in start state. [PR1460555](#)
- 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 (aggregated Ethernet1,

50 percent line rate 4 streams bi-directional => 118-Gbps combined traffic rate) both hosted on a single Packet Forwarding Engine instruction of an MPC11E line card, small varying packet drops occurs at 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 **HALP-trinity_nh_dynamic_mcast_add_irb_topo:3520 snooping-error: invalid IRB topo/ IRB ifl zero in I2 nh 40495 add IRB.** [PR1472222](#)
- Line-card errors found at **RT-HAL,rt_mesh_group_delete_check,1599: Deletion of a non-existent 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 a unified ISSU, when you enable or disable and activate or deactivate the interface. [PR1493723](#)
- 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](#)
- With subscriber services configuration and distributed IGMP processing enabled for subscribers, it is possible that the line card can occasionally crash due to npc process core file generation. A line-card reboot is required to recover. This issue is seen outside of subscriber services or even with subscriber services if distributed IGMP is not enabled. [PR1534542](#)
- On MX480 routers, IPv6 VRRP sessions are not established when duplicate address detection (DAD) is enabled. [PR1534835](#)
- On certain Junos platforms with dual Routing Engines (platforms capable of installing Junos packages with name format as "junos*install"), BGP replication might fail to start under GRES/NSR setup after a crash on backup Routing Engine. NSR starts unreplicating the socket since backup Routing Engine is no longer present. Massive unreplicated request leads to memory buffer getting full with multiple BGP

sessions (for example, 20 BGP peers). Hence BGP unreplicated request returned with an error. Besides, the kernel is left with stale data. It does not allow the Juniper Socket Replication (JSR) when backup Routing Engine comes up due to the stale data. BGP-NSR is broke under the conditions. Traffic outage will be observed after performing GRES. [PR1552603](#)

Routing Policy and Firewall Filters

- Routing policy actions failed to configure **neighbor-sets** and **tag-sets**. [PR1491795](#)

Routing Protocols

- While interoperateing 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](#)
- When the device is configured with link-node protection and labeled-bgp, ukern memory leak and FPC core file generation might occur. [PR1366823](#)
- BFD session flaps during unified ISSU in the MPC7E line card. [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 the inet.3 table irrespective of 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 scaled scenario in a very short period of time dedicated PIM router receives more than 2500 PIM hello packets from the new neighbors followed by PIM Join packets for the same multicast group, some of this joins might not be processed from the first attempt. [PR1500125](#)
- On devices with Next Generation Routing Engine and SCBE2 (Enhanced Switch Control Board), when BFD authentication for BGP is enabled, the BFD might flap after a Next Generation Routing Engine switchover. After the flap, the device does a self-recovery. [PR1522261](#)
- When VRRP virtual address is configured and used to set up a BGP session with the remote side, under rare timing conditions, BGP peer establishment might get rejected repetitively. [PR1523075](#)
- On MX Series with MPCs/MICs based Virtual Chassis, when there are multicast tunneled packets being received, which come into the Virtual Chassis Ports (VCP) and then pop out of the tunnel, if the VCP ports and the interfaces where multicast packets enter/leave the router are located on the same Packet

Forwarding Engine, it might fail in sending multicast traffic to downstream receiver due to this issue. [PR1555518](#)

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](#)

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](#)
- The MPC10E-15C-MRATE next-generation MPVPN ingress replication flushing out is not proper when in egress the ingress replication configuration is deactivated. [PR1475834](#)
- In a multicast VPN RPT-SPT mode with both locally and remotely connected receivers, the multicast forwarding entry related to the actual downstream interfaces might not be properly updated. The issue might happen when multicast forwarding entry is created by locally connected PIM receiver, followed by remote receiver and after local receiver decides to prune its membership in that multicast group. [PR1546739](#)

SEE ALSO

[What's New | 78](#)

[What's Changed | 108](#)

[Known Limitations | 114](#)

[Resolved Issues | 130](#)

[Documentation Updates | 165](#)

[Migration, Upgrade, and Downgrade Instructions | 166](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.1R2 | 130](#)
- [Resolved Issues: 20.1R1 | 149](#)

This section lists the issues fixed in Junos OS Release 20.1R2 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: 20.1R2

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_icode_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, dhcpcd 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](#)
- GnmJuniperTelemetryHeader 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 multiple firmwares for different hardware. [PR1524338](#)
- Commit is successful while deactivating CBO 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 DHCPLEASEUNASSIGNED. [PR1512239](#)
- The **show dhcpv6 relay statistics** command must display DHCPV6LEASEQUERY_REPLY instead of DHCPV6LEASEQUERY_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 lcemd, 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 Penta 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 (0x00000000088488c 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 mcsnooopd 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 **agentRemotId** 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 | 78](#)
- [What's Changed | 108](#)
- [Known Limitations | 114](#)
- [Open Issues | 117](#)
- [Documentation Updates | 165](#)
- [Migration, Upgrade, and Downgrade Instructions | 166](#)

Documentation Updates

IN THIS SECTION

- [Dynamic Host Configuration Protocol \(DHCP\) | 165](#)

This section lists the errata and changes in Junos OS Release 20.1R2 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

- [What's New | 78](#)
- [What's Changed | 108](#)
- [Known Limitations | 114](#)
- [Open Issues | 117](#)
- [Resolved Issues | 130](#)
- [Migration, Upgrade, and Downgrade Instructions | 166](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.1R2 | 167](#)
- [Procedure to Upgrade to FreeBSD 11.x based Junos OS | 167](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 170](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 171](#)
- [Upgrading a Router with Redundant Routing Engines | 172](#)
- [Downgrading from Release 20.1R2 | 172](#)

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

[What's New | 78](#)

[What's Changed | 108](#)

[Known Limitations | 114](#)

[Resolved Issues | 130](#)

[Open Issues | 117](#)

Documentation Updates | 165

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [What's New | 173](#)
- [What's Changed | 175](#)
- [Known Limitations | 176](#)
- [Open Issues | 176](#)
- [Resolved Issues | 177](#)
- [Documentation Updates | 180](#)
- [Migration, Upgrade, and Downgrade Instructions | 181](#)

These release notes accompany Junos OS Release 20.1R2 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

- [What's New in Release 20.1R2 | 174](#)
- [What's New in Release 20.1R1 | 174](#)

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 Series, 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 | 175](#)
-
- [Known Limitations | 176](#)
-
- [Open Issues | 176](#)
-
- [Resolved Issues | 177](#)
-
- [Documentation Updates | 180](#)
-
- [Migration, Upgrade, and Downgrade Instructions | 181](#)

What's Changed

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

SEE ALSO

-
- [What's New | 175](#)
-
- [Known Limitations | 176](#)

[Open Issues | 176](#)

[Resolved Issues | 177](#)

[Documentation Updates | 180](#)

[Migration, Upgrade, and Downgrade Instructions | 181](#)

Known Limitations

There are no known limitations for NFX Series devices in Junos OS Release 20.1R2.

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 | 173](#)

[What's Changed | 175](#)

[Open Issues | 176](#)

[Resolved Issues | 177](#)

[Documentation Updates | 180](#)

[Migration, Upgrade, and Downgrade Instructions | 181](#)

Open Issues

IN THIS SECTION

- [Interfaces | 177](#)

- [Platform and Infrastructure | 177](#)

- [Virtual Network Functions \(VNFs\) | 177](#)

Learn about open issues in Junos OS Release 20.1R2 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.

Interfaces

- On NFX350 devices, the **clear interface statistics all** command takes a longer time to execute. [PR1475804](#)

Platform and Infrastructure

- On NFX150 devices, an srxpfe core file is observed while testing the ADSL interface. [PR1485384](#)
- If you plug an unsupported SFP-T transceiver into an NFX150 device and reboot the device, the FPC1 WAN port does not come online. [PR1411851](#)
- On NFX350 devices, traffic drop is seen with fragmented traffic, and the log reports **FLOW_REASSEMBLE_FAIL**. [PR1475023](#)

Virtual Network Functions (VNFs)

- 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 | 173](#)
- [What's Changed | 175](#)
- [Known Limitations | 176](#)
- [Resolved Issues | 177](#)
- [Documentation Updates | 180](#)
- [Migration, Upgrade, and Downgrade Instructions | 181](#)

Resolved Issues

IN THIS SECTION

-  [Resolved Issues: 20.1R2 | 178](#)
-  [Resolved Issues: 20.1R1 | 179](#)

Learn which issues were resolved in the Junos OS Release 20.1R1 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: 20.1R2

Application Security

- AppQoE is sending active probe packets for the deleted **active-probe-params**. [PR1492208](#)

Interfaces

- On NFX250 NextGen devices, 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 **d10** 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 | 173](#)
-
- [What's Changed | 175](#)
-
- [Known Limitations | 176](#)
-
- [Open Issues | 176](#)
-
- [Documentation Updates | 180](#)
-
- [Migration, Upgrade, and Downgrade Instructions | 181](#)

Documentation Updates

There are no errata or changes in Junos OS Release 20.1R2 documentation for NFX Series.

SEE ALSO

-
- [What's New | 173](#)
-
- [What's Changed | 175](#)
-
- [Known Limitations | 176](#)
-
- [Open Issues | 176](#)

[Resolved Issues | 177](#)[Migration, Upgrade, and Downgrade Instructions | 181](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 181](#)
- [Basic Procedure for Upgrading to Release 20.1 | 181](#)

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 20.1R2:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the **Software** tab.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the Download Software page.
5. In the Install Package section of the Software tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the device or to your internal software distribution site.
10. Install the new package on the device.

SEE ALSO

- [What's New | 173](#)
- [What's Changed | 175](#)
- [Known Limitations | 176](#)
- [Open Issues | 176](#)
- [Resolved Issues | 177](#)
- [Documentation Updates | 180](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- [What's New | 184](#)
- [What's Changed | 191](#)
- [Known Limitations | 192](#)
- [Open Issues | 194](#)
- [Resolved Issues | 197](#)
- [Documentation Updates | 201](#)
- [Migration, Upgrade, and Downgrade Instructions | 202](#)

These release notes accompany Junos OS Release 20.1R2 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

- [What's New in 20.1R2 | 184](#)
- [What's New in 20.1R1 | 184](#)

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

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 device 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 (2⁰ to 2⁶).

```
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 | 191](#)

[Known Limitations | 192](#)

[Open Issues | 194](#)

[Resolved Issues | 197](#)

[Documentation Updates | 201](#)

[Migration, Upgrade, and Downgrade Instructions | 202](#)

What's Changed

IN THIS SECTION

- [What's Changed in 20.1R2 Release | 191](#)

See what changed in this release for PTX Series routers.

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](#).]

SEE ALSO

[What's New | 184](#)

[Known Limitations | 192](#)

[Open Issues | 194](#)

[Resolved Issues | 197](#)

[Documentation Updates | 201](#)

[Migration, Upgrade, and Downgrade Instructions | 202](#)

Known Limitations

IN THIS SECTION

 [General Routing | 193](#)

 [MPLS | 194](#)

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](#)
- 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 | 184](#)

[What's Changed | 191](#)

[Open Issues | 194](#)

[Resolved Issues | 197](#)

[Documentation Updates | 201](#)

[Migration, Upgrade, and Downgrade Instructions | 202](#)

Open Issues

IN THIS SECTION

- [General Routing | 195](#)
- [Infrastructure | 196](#)
- [MPLS | 196](#)
- [Routing Protocols | 196](#)

Learn about open issues 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

- When CFP2-DCO-T-WDM-1 is plugged in to a PTX Series PIC, after FPC restarts, the carrier frequency offset TCA is raised even when TCA is not enabled. [PR1301471](#)
- Alarm action does not work for minor errors after changing the threshold to 1. [PR1345154](#)
- Traffic loss is greater than 50 ms (in order of 200 to 300 ms) for IP routes pointing to a unilist of composites with indirect next hops during a link down scenario. In this case, the Packet Forwarding Engine does not do the local repair and will wait for the rpd to install the new next hops. [PR1383965](#)
- On the PTX3000 routers, the firewall counter for lo0 does not increment. [PR1420560](#)
- 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. The detached link is invalidated from the Packet Forwarding Engine AE-Forwarding table as expected. However, if the device is rebooted with this state, all the member links are enabled in the Packet Forwarding Engine AE-Forwarding table irrespective of LACP states, which results in traffic drop. [PR1423707](#)
- The em2 interface configuration is causing the FPC to crash during initialization, and the FPC does not come online. After deleting the em2 configuration and restarting the router, the FPC comes online. [PR1429212](#)
- Memory leak observed while running the longevity check. [PR1438358](#)
- When users configure the best destination network with **dyn-tunnel-attribute-policy** and preference, the tunnel is not migrated from the old destination network. [PR1462805](#)
- 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](#)
- On PTX10016 routers, after device reboot, the FPC takes a long time to come up and hence MKA sessions establishment is delayed. The error message **Frame 08: sp = 0x48d222b8, pc = 0x10fad3bc , blaze fpc2 SCED: Thread 59 (PFE Manager) ran for 2177 ms without yielding** is observed. [PR1477585](#)
- The Layer 2 VPN with asynchronous-notification configured might flap when the link goes up between the PE and CE devices. After Layer 2 VPN flaps, the interfaces with asynchronous-notification might show **- Inf dBm** laser output power even if the Layer 2 VPN is in the up state. [PR1486181](#)
- In strict-priority scheduling mode, medium-high and medium-low are seen to be operating on the same priority. [PR1490505](#)
- Junos telemetry interface GRPC multisensors do not work as expected. [PR1492282](#)

Infrastructure

- 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, and 16.2R1). [PR1222105](#)

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 and then 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 cached value of traffic statistics, which might not be equal to the current value. [PR1445705](#)
- With NSR enabled, the current BGP design supports 3000 BGP IPv6 peers or 8000 BGP IPv4 peers. When you 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 | 184](#)

[What's Changed | 191](#)

[Known Limitations | 192](#)

[Resolved Issues | 197](#)

[Documentation Updates | 201](#)

[Migration, Upgrade, and Downgrade Instructions | 202](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.1R2 | 197](#)
- [Resolved Issues: 20.1R1 | 199](#)

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: 20.1R2

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: ieee8021CfmStackServiceSelectorType.** [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 I2cpd 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 asynchronized 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 | 184](#)

[What's Changed | 191](#)

[Known Limitations | 192](#)

[Open Issues | 194](#)

[Documentation Updates | 201](#)

[Migration, Upgrade, and Downgrade Instructions | 202](#)

Documentation Updates

IN THIS SECTION

- [Dynamic Host Configuration Protocol \(DHCP\) | 201](#)

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

[What's New | 184](#)

[What's Changed | 191](#)

[Known Limitations | 192](#)

[Open Issues | 194](#)

[Resolved Issues | 197](#)

[Migration, Upgrade, and Downgrade Instructions | 202](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.1 | 202](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 205](#)
- [Upgrading a Router with Redundant Routing Engines | 205](#)

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 | 184](#)
- [What's Changed | 191](#)
- [Known Limitations | 192](#)

- [Open Issues | 194](#)
- [Resolved Issues | 197](#)
- [Documentation Updates | 201](#)

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- [What's New | 206](#)
- [What's Changed | 212](#)
- [Known Limitations | 216](#)
- [Open Issues | 219](#)
- [Resolved Issues | 224](#)
- [Documentation Updates | 237](#)
- [Migration, Upgrade, and Downgrade Instructions | 238](#)

These release notes accompany Junos OS Release 20.1R2 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

- [What's New in Release 20.1R2 | 207](#)
- [What's New in Release 20.1R1 | 208](#)

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 20.1R2: 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 20.1R2

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 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](#).]

- **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 device 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`, `smpete`, or `aes67-smpete` 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 | 212](#)

[Known Limitations | 216](#)

[Open Issues | 219](#)

[Resolved Issues | 224](#)

[Documentation Updates | 237](#)

[Migration, Upgrade, and Downgrade Instructions | 238](#)

What's Changed

IN THIS SECTION

- [What's Changed in 20.1R2 | 213](#)
- [What's Changed in 20.1R1 | 214](#)

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

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-Mbps 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 | 206](#)

[Known Limitations | 216](#)

[Open Issues | 219](#)

[Resolved Issues | 224](#)

[Documentation Updates | 237](#)

[Migration, Upgrade, and Downgrade Instructions | 238](#)

Known Limitations

IN THIS SECTION

- [Class of Service \(CoS\) | 217](#)
- [EVPN | 217](#)
- [Infrastructure | 217](#)
- [Layer 2 Ethernet Services | 217](#)
- [Layer 2 Features | 217](#)
- [Platform and Infrastructure | 217](#)

Learn about known limitations in Junos OS Release 20.1R2 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 QFX5100 and EX4600 platforms, due to a major third-party SDK upgrade in Junos OS Release 20.1R1 (from SDK 6.3.7 to 6.5.16), ISSU is not supported from any earlier releases to Junos OS Release 20.1 (image : jinstall-qfx-5-*). [PR1479439](#)

EVPN

- BUM traffic loop happens when a core interface is isolated in a short time period. The loop prevention feature does not work because the CE interface is not flapped with core interface isolation. The Routing Engine sets CE interface down after the BGP hold time has expired, and the Routing Engine sets the CE interface to up after the BGP session is established. The default BGP hold time is 90 seconds. It needs 90 seconds to flap the CE interface. The BUM traffic can be looped till the local bias filter is restored. [PR1492784](#)

Infrastructure

- If Junos OS experiences a file system-related panic, such as 'dup alloc', recovery through the OAM shell might be needed. From the OAM shell, run 'fsck' on the root volume until it is marked clean. Only at this point is it safe to reboot to the normal volume. [PR1444941](#)

Layer 2 Ethernet Services

- In an EVPN multihomed active-active scenario, when LACP is enabled on PE-CE child member links, the LACP force-up feature should not be enabled in conjunction with the EVPN core isolation feature (enabled by default) because it is currently not supported in this scenario as these two features are contradictory in terms of action they take. [PR1461581](#)

Layer 2 Features

- On QFX5000 platforms, you might see the **pools exhausted for Table:EGR_DVP_ATTRIBUTE** error message when statistics requests exceeded the supported scale because of the limited pool resources used for statistics collection on the hardware. There is no functional impact except for the statistics collection for some hardware counters for which flex counter allocation failed for the time, the limit is exceeded. The statistics counters start functioning normally without manual change when the pool comes back to the normal limit. [PR1479826](#)

Platform and Infrastructure

- Downgrade from a TVP image to a non-TVP image is not supported. However, upgrade from a non-TVP image to a TVP image is supported. [PR1345848](#)
- The 100-Gigabit Ethernet interface goes down after you configure and delete the Ethernet loopback configuration. [PR1353734](#)
- On the QFX10000 line of devices, if an analyzer is configured to mirror traffic of an input aggregated Ethernet interface and a new member is added to the same aggregated Ethernet interface, then the analyzer might not provide sample packets that flow through a newly added child interface. [PR1417694](#)
- During software validation, Junos OS mounts the new image and validates the configuration against the new image. As QFX5000 and QFX10000 platforms are already mounting the maximum four disks during normal execution, Junos OS cannot mount the extra disk for this purpose. Thus QFX currently does not support configuration validation during upgrade on QFX5000, so the syntax error appears when the image installation is triggered with "validation". [PR1421378](#)
- Due to additional hi-gig header, 100% throughput cannot be achieved when packets are forwarded through Virtual Chassis ports. For 64-byte packets, throughput is ~91% and for 1024-byte packets the throughput is ~99%. [PR1453709](#)
- The interface with a new 10G SFP-T on an enable/disable has for a second displays media as Fiber instead of Copper. The present code has brought in 10G copper functionality and this is done without the capability addition to the XE physical interfaces. [PR1467509](#)
- Under the **set chassis nssu upgrade-group <xx> fpc** configuration command if multiple FPCs are included; then the NSSU fails. The recommendation is to have a single FPC per upgrade group for all the FPCs in the VCF setup. This only affects VCF setup. [PR1473624](#)
- Convergence delay for link-protected MPLS LSP is more than 50 ms. [PR1478584](#)
- [evpn_vxlan] [evpn_instance] Observing 100% L2 MAC scaling traffic loss in QFX10002-60C platform after loading evpn-vxlan collapsed profile configurations. [PR1489753](#)
- On QFX5100, NSSU from an older Junos OS release with QFX5100-VC SDK 6.3.x to a new Junos OS release with Broadcom SDK 6.5.x may not work. As a workaround, a normal upgrade from older release to the new release can be done. [PR1496765](#)

SEE ALSO

[What's New | 206](#)

[What's Changed | 212](#)

[Open Issues | 219](#)

[Resolved Issues | 224](#)

[Documentation Updates | 237](#)

[Migration, Upgrade, and Downgrade Instructions | 238](#)

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 220](#)
- [EVPN | 220](#)
- [High Availability \(HA\) and Resiliency | 220](#)
- [Interfaces and Chassis | 220](#)
- [Layer 2 Ethernet Services | 221](#)
- [Layer 2 Features | 221](#)
- [Platform and Infrastructure | 221](#)
- [Routing Protocols | 223](#)
- [Virtual Chassis | 224](#)

Learn about open issues in Junos OS Release 20.1R2 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)

- When CPU Q cells (memory) are exhausted, even though the incoming packet rate is less than the allowed bandwidth, you observe DDoS violations. [PR1381775](#)

EVPN

- Bum traffic loop happen when core interface is isolated in short time period. Loop prevention feature does not work because CE interface is not flapped with core interface isolation. Routing Engine sets CE interface down after BGP hold time expired and Routing Engine sets CE interface to up after BGP session established. Default BGP hold time is 90 seconds. It needs 90 seconds to flap CE interface. Bum traffic can be looped till local bias filter is restored. [PR1492784](#)

High Availability (HA) and Resiliency

- An issue was reported for a customer with a Flush Cache issue on the same platform. As it was root-caused to a reliable SSD Disk I/O change to be made for this platform, this caused the added delay observed in the reported issue. [PR1511607](#)

Interfaces and Chassis

- Multicast traffic can be flooded for 15 to 20 seconds to both MC-LAG peers, after the following sequence of steps:
 1. Disable or enable ICL.
 2. Reboot one of MC-LAG peers.
 3. Disable or enable a member link of ICL. This results in no traffic loss, and one of the MC-LAG nodes processes duplicate packets during this time period. [PR1422473](#)
- Flooding of ARP reply unicast packets is seen as a result of an ARP request sent for the device's VRRP MAC address. The ARP reply, which is flooded in the VLAN by the device, has the correct DMAC of the originator of the ARP request. In other words, the ARP reply is flooded but with the correct unicast DMAC. The ARP reply is not broadcast. It is independent of MC-LAG and VRRP scenarios. [PR1454764](#)

Layer 2 Ethernet Services

- If forward-only is set within dhcp-reply in a Juniper Networks 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](#)

Layer 2 Features

- In case of QFX5000 Virtual Chassis/VCF setups, when IGMP snooping is enabled, multicast traffic is forwarded based on IGMP joins/reports. But when the IGMP report times out, traffic should be dropped, instead it is flooded in the VLAN. This happens only in case of QFX5000 Virtual Chassis/VCF, this issue is not seen on stand-alone QFX5000. [PR1431893](#)
- On QFX5000 platforms, you might see the **pools exhausted for Table:EGR_DVP_ATTRIBUTE** error message when statistics requests exceeded the supported scale because of the limited pool resources used for statistics collection on the hardware. There is no functional impact except for statistics collection for some hardware counters for which flex counter allocation failed for the time, the limit is exceeded. The statistics counters start functioning normally without manual change when the pool comes back to normal limit. [PR1479826](#)
- On QFX5100 switches, FXPC CPU utilization is increased due to high number of active ports after third-party SDK upgrade to 6.5.x from 5.3.x. [PR1480132](#)

Platform and Infrastructure

- On all platforms running Junos OS that support EVPN-MPLS/EVPN-VXLAN, when an existing ESI interface flaps or is added newly to the configuration, sometimes DF (Designated Forwarder) election happens before the local bias feature is enabled and during this time, existing BUM (Broadcast, Unknown unicast, Multicast) traffic might be looped for a short time duration (less than several seconds). [PR1493650](#)
- Port LEDs on the QFX5100 do not work. If a device connects to a port on the QFX5100, the port LED stays unlit. [PR1317750](#)
- QFX10000:Source MAC and TTL values are not updated for routed multicast packets in EVPN-VXLAN. [PR1346894](#)
- QFX10000 platform drops the Aruba wireless access point (AP) heartbeat packets. As a result, the Aruba wireless AP cannot work. [PR1352805](#)
- USB upgrade of NOS image is not supported. [PR1373900](#)
- On Junos OS Release 18.4R1 branch, intermittent traffic loss is observed with RTG streams while flapping the RTG primary interface. [PR1388082](#)
- On QFX5110 and QFX5120 platforms, either unicast RPF in strict mode or ICMP redirect does not work properly. [PR1417546](#)

- 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 occurs when a new member link is added where the other end of that link is not yet configured with LACP. [PR1423707](#)
- When you restart the routing process, if the system is configured with EVPN service, memory of the Layer 2 learning daemon increases by 4000 when you use **show system processes extensive | match l2ald**. [PR1435561](#)
- Unified ISSU is not supported on QFX5200 switches and fails from Junos OS Release 17.2X75-D43.2 to some target versions. Also, dcpfe crash might be seen. [PR1438690](#)
- QFX5000 platforms support the port qualifier. This will install two entries in Packet Forwarding Engine, one with source-port and second one with destination-port with value as what is specified under the port stanza. [PR1440980](#)
- On QFX10000 platforms, in an EVPN-VXLAN (spine-leaf) scenario, the QFX10000 spine switches are configured with the VXLAN Layer 3 gateway (utilizing the virtual gateway) on an IRB interface. If you enable and then subsequently remove the VXLAN Layer 3 gateway on this IRB interface on one or some of these spine switches, traffic drop might be observed. As a workaround, configure all virtual gateways with unique IPv4 or IPv6 MAC addresses. [PR1446291](#)
- On QFX5000, triggering NSSU on a Virtual Chassis will print ISSU logs as NSSU uses the same state machine as ISSU. There is no functional impact due to this behavior. [PR1451375](#)
- Whenever any member in a remote Switched Port Analyzer (RSPAN) VLAN is removed from that VLAN, you must reconfigure the analyzer session for that RSPAN VLAN. [PR1452459](#)
- In overall commit time, the evaluation of mustd constraints is taking 2 seconds more than usual. This is because the persist-group-inheritance feature has been made a default feature in the latest Junos OS releases. 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. Because issue is seen only with QFX Series platforms or other low end devices, we can release note commit time degradation behavior only for QFX platforms. [PR1457939](#)
- The VXLAN VNI (multicast learning) scaling on QFX5110 traffic issue is seen from VXLAN tunnel to L2 interface. [PR1462548](#)
- The output of the **show chassis environment** command can be seen from backup members as well. The issue is common to all QFX Series platforms. [PR1474520](#)
- Interfaces are not detected on some of the ports when you swap the 25-gigabit SFP transceiver and insert a 10-gigabit transceiver. [PR1475574](#)
- On QFX5220 platforms with Junos OS Evolved images, when the lo0 firewall filter (inet/inet6) is used, Layer 3 forwarding traffic might be discarded by the lo0 filter. [PR1475620](#)

- **LBCM-L2,pfe_shm_vrf_hw_token_map_add(),4987:MHOP pfe_shm_vrf_hw_token_map_add parameters are wrong** error observed after loading baseconfig. [PR1480149](#)
- Instead of disabling ptpt0 port, you can delete the ptpt0 interface to achieve the same result. [PR1487505](#)
- The **delete openconfig-routing-policy [prefix-set-name]** CLI command did not work, and it failed with error **OC Deletion did not work properly.** [PR1492561](#)
- On QFX5100, NSSU from an older Junos OS release with QFX5100-VC SDK 6.3.x to a new Junos OS release with Broadcom SDK 6.5.x may not work. As a workaround, normal upgrade from the older release to the new release can be done. [PR1496765](#)
- Users cannot subscribe to any path that ends with "key". [PR1553534](#)
- Domain and VLAN parameters are not as expected while verifying **show dot1x interface detail**. Expected VoIP but received Data. [PR1553596](#)
- Disruptive switchover (no GRES or NSR configured) can lead to stale PPM (Periodic Packet Management) entries programmed on the new primary Routing Engine, if both GRES and NSR are activated after a disruptive switchover and then a GRES is performed, BFD sessions might flap continuously. [PR1518106](#)
- On a device running Junos OS, in an EVPN scenario, if an interface is in transition status and an ARP request is received on the interface from a host, the device may send out re-arp toward the host, and the host responds to this re-arp and the Junos device sends another re-arp in response to this arp, reply from host. This goes on forever, causing a high rate of ARP packets until the interface comes up. This issue is also applicable to ND/NS in an IPv6 environment. [PR1534796](#)

Routing Protocols

- DCPFE core-files after watchdog trigger caused by the failed MAC deletion notification. The following repeated messages before the core-dump can be an evidence of the problem:
BRCM_SALM:brcm_salm_periodic_clear_pending(),125: Failed to delete Pendingentries for unit = 0, tgid = 1, err code = -9. The fix is present in 17.3R4, 17.4R2, 18.1R3, 18.2R1, 17.3R2-S4, 17.3R3-S1, 15.1X53-D235 and later releases. [PR1371092](#)
- On QFX5100 Virtual Chassis or Virtual Chassis Fabric, when the mini-PDT-base configuration is issued, the following error message is seen in the hardware: **BRCM_NH-,brcm_nh_bdvlans_ucast_uninstall(), 128:I3 nh 6594 unintsall failed**. There is no functionality impact because of this error message. [PR1407175](#)
- On QFX10000 platforms, if multiple sub-interfaces of the same Aggregated Ethernet (AE) interface are belonged to different routing instances, and these sub-interfaces are configured with the same IP address and configured with separate BFD (Bidirectional Forwarding Detection) sessions, the remaining BFD sessions will flap continuously if one of these BFD sessions is deleted. [PR1516556](#)

Virtual Chassis

- 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` are seen. This is a false positive error message. [PR1276970](#)
- On QFX5000 Virtual Chassis, DDoS violations on backup are not reported to the Routing Engine. [PR1490552](#)

SEE ALSO

[What's New | 206](#)

[What's Changed | 212](#)

[Known Limitations | 216](#)

[Resolved Issues | 224](#)

[Documentation Updates | 237](#)

[Migration, Upgrade, and Downgrade Instructions | 238](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.1R2 | 224](#)
- [Resolved Issues: 20.1R1 | 230](#)

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: 20.1R2

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 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 I2cpd process might crash if the ERP configuration is added or removed, and the I2cpd 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](#)
- The 100GbE 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_lrnct(),1379:jprds_slu_sal_update_lrnct 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 classifier, 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 Boxes (-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 asynchronized 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 mcsnooopd 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 | 206](#)
- [What's Changed | 212](#)
- [Known Limitations | 216](#)
- [Open Issues | 219](#)
- [Documentation Updates | 237](#)
- [Migration, Upgrade, and Downgrade Instructions | 238](#)

Documentation Updates

IN THIS SECTION

- [Dynamic Host Configuration Protocol \(DHCP\) | 237](#)

This section lists the errata and changes in Junos OS Release 20.1R2 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 | 206](#)
- [What's Changed | 212](#)
- [Known Limitations | 216](#)
- [Open Issues | 219](#)
- [Resolved Issues | 224](#)
- [Migration, Upgrade, and Downgrade Instructions | 238](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrading Software on QFX Series Switches | 238](#)
- [Installing the Software on QFX10002-60C Switches | 241](#)
- [Installing the Software on QFX10002 Switches | 241](#)
- [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 | 242](#)
- [Installing the Software on QFX10008 and QFX10016 Switches | 244](#)
- [Performing a Unified ISSU | 248](#)
- [Preparing the Switch for Software Installation | 249](#)
- [Upgrading the Software Using Unified ISSU | 249](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 251](#)

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-x86-64-20.1R2.9.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-x86-64-20.1R2.9.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-x86-64-20.1R2.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-x86-64-20.1R2.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-x86-64-20.1R2.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-x86-64-20.1R2.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 249](#)
- [Upgrading the Software Using Unified ISSU on page 249](#)

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-x86-64-20.1R2.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 | 206](#)

[What's Changed | 212](#)

[Known Limitations | 216](#)

[Open Issues | 219](#)

[Resolved Issues | 224](#)

[Documentation Updates | 237](#)

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [What's New | 253](#)
- [What's Changed | 260](#)
- [Known Limitations | 264](#)
- [Open Issues | 266](#)
- [Resolved Issues | 268](#)
- [Documentation Updates | 278](#)
- [Migration, Upgrade, and Downgrade Instructions | 279](#)

These release notes accompany Junos OS Release 20.1R2 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

- [Release 20.1R2 New and Changed Features | 253](#)
- [Release 20.1R1 New and Changed Features | 253](#)

Learn about new features introduced in the Junos OS main and maintenance releases for 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 Series, 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 | 260](#)

[Known Limitations | 264](#)

[Open Issues | 266](#)

[Resolved Issues | 268](#)

[Documentation Updates | 278](#)

[Migration, Upgrade, and Downgrade Instructions | 279](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.1R2 | 261](#)
- [What's Changed in Release 20.1R1 | 262](#)

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

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

[What's New | 253](#)

[Known Limitations | 264](#)

[Open Issues | 266](#)

[Resolved Issues | 268](#)

[Documentation Updates | 278](#)

[Migration, Upgrade, and Downgrade Instructions | 279](#)

Known Limitations

IN THIS SECTION

● [J-Web | 265](#)

● [Platform and Infrastructure | 265](#)

● [VPNs | 265](#)

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.

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](#)

Platform and Infrastructure

- 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](#)

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](#)

SEE ALSO

[What's New | 253](#)

[What's Changed | 260](#)

[Open Issues | 266](#)

[Resolved Issues | 268](#)

[Documentation Updates | 278](#)

[Migration, Upgrade, and Downgrade Instructions | 279](#)

Open Issues

IN THIS SECTION

- [ATP Cloud | 266](#)
- [Flow-Based and Packet-Based Processing | 266](#)
- [J-Web | 267](#)
- [Routing Policy and Firewall Filters | 267](#)
- [VPNs | 267](#)

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.

ATP Cloud

- CLI autocomplete is now available for both SeclIntel and advanced anti-malware products. [PR1487419](#)

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](#)
- 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](#)

- The MAC table is null in Layer 2 mode after one pass-through session is created successfully. [PR1528286](#)

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](#)
- The firewall Web authentication graphics have been updated. [PR1482433](#)

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](#)
- An issue was discovered within unified policies that affected the url-category match condition that can cause it to over-match and apply to more traffic than it should. [PR1546120](#)

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](#)
- The command show security pki local-certificate logical-system all is not showing any output. [PR1414628](#)
- 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 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](#)
- IPsec VPN flaps if more than 500 IPsec VPN tunnels are connected for the first time. [PR1455951](#)
- In an IPsec VPN scenario on the SRX5000 line of devices, 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 amount 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 struck after RG1 failover. [PR1477184](#)
- During 10,000 tunnel ramp-up, sometimes, IKED generates a core file. [PR1479548](#)
- On SRX5000 line of devices with SPC3, the tunnel is not brought down immediately after disabling the interface of the peer device with DPD always-send configured on site-2-site route-based VPN. [PR1480905](#)
- On SRX5000 line of devices, with extended-sequence-number (ESN) configured, the IPsec tunnel might be re-established. This issue could be self-recovered, and traffic loss happens during IPsec tunnel flapping. [PR1488087](#)
- Unexpected extra characters NL were seen with PyEZ XML outputs. This caused issues while writing op-scripts. However, with normalize=True in PyEz script, we 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 was not supporting simultaneous IKE negotiation in Junos OS Releases 19.2, 19.3, 19.4, and 20.1. [PR1497297](#)
- IPsec traffic might get dropped after RG0 failover. [PR1522931](#)

SEE ALSO

[What's New | 253](#)

[What's Changed | 260](#)

[Known Limitations | 264](#)

[Resolved Issues | 268](#)

[Documentation Updates | 278](#)

[Migration, Upgrade, and Downgrade Instructions | 279](#)

Resolved Issues

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: 20.1R2

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 | 253](#)

[What's Changed | 260](#)

[Known Limitations | 264](#)

[Open Issues | 266](#)

[Documentation Updates | 278](#)

[Migration, Upgrade, and Downgrade Instructions | 279](#)

Documentation Updates

IN THIS SECTION

- [Dynamic Host Configuration Protocol \(DHCP\) | 279](#)

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 | 253](#)

[What's Changed | 260](#)

[Known Limitations | 264](#)

[Open Issues | 266](#)

[Resolved Issues | 268](#)

[Migration, Upgrade, and Downgrade Instructions | 279](#)

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 | 253](#)
- [What's Changed | 260](#)
- [Known Limitations | 264](#)
- [Open Issues | 266](#)
- [Resolved Issues | 268](#)
- [Documentation Updates | 278](#)

Junos OS Release Notes for vMX

IN THIS SECTION

- [What's New | 281](#)
- [What's Changed | 281](#)
- [Known Limitations | 281](#)
- [Open Issues | 281](#)
- [Resolved Issues | 281](#)
- [Licensing | 282](#)
- [Upgrade Instructions | 283](#)

These release notes accompany Junos OS Release 20.1R2 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

There are no new features for vMX in Junos OS Release 20.1R2.

What's Changed

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 20.1R2.

Open Issues

There are no open issues for vMX in Junos OS Release 20.1R2.

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.

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

- [What's New | 283](#)
- [What's Changed | 284](#)
- [Known Limitations | 285](#)
- [Open Issues | 286](#)
- [Resolved Issues | 287](#)
- [Migration, Upgrade, and Downgrade Instructions | 289](#)

These release notes accompany Junos OS Release 20.1R2 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

- [What's New in Release 20.1R2 | 284](#)

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

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

- [What's Changed in Release 20.1R2 | 285](#)

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

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

- [Flow-Based and Packet-Based Processing | 286](#)
- [J-Web | 286](#)
- [User Access and Authentication | 286](#)

Learn about known limitations in Junos OS Release 20.1R2 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. These algorithms have better performance on vSRX than CBC encryption algorithms, such as aes-128-cbc. [PR1444022](#)

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](#)

User Access and Authentication

- 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 address 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 the underscore "_" in the certificate ID field while creating keypairs. This is a limitation from Azure KeyVault. [PR1475254](#)

Open Issues

IN THIS SECTION

- [Intrusion Detection and Prevention \(IDP\) | 287](#)

Learn about open issues in Junos OS Release 20.1R2 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: 20.1R2

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

- [Upgrading Software Packages | 290](#)
- [Validating the OVA Image | 295](#)

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 20.1R2 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 20.1R2 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 20.1R2 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-x86-64-20.1R2-2020-9-10.0_RELEASE_20.1R2_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-x86-64-20.1R2-2020-9-10.0_RELEASE_20.1R2_THROTTLE.tgz
no-copy no-validate reboot
Verified junos-vsrx-x86-64-20.1R2-2020-9-10.0_RELEASE_20.1R2_THROTTLE signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING: This package will load JUNOS 20.1R2 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-20.1R2-2020-9-10.0_RELEASE_20.1R2_THROTTLE.tgz
Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31

```

```
junos-srx-mr-vsrx-20.1R2-2020-9-10.0_RELEASE_20.1R2_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31
junos-srx-mr-vsrx-20.1R2-2020-9-10.0_RELEASE_20.1R2_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-20.1R2-2020-9-10.0_RELEASE_20.1R2_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-20.1R2-2020-9-10.0_RELEASE_20.1R2_THROTTLE-linux.tgz
...
upgrade_platform: Input package
/var/tmp/junos-srx-mr-vsrx-20.1R2-2020-9-10.0_RELEASE_20.1R2_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-20.1R2-2020-9-10.0_RELEASE_20.1R2_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-20.1R2-2020-9-10.0_RELEASE_20.1R2_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 20.1R2 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 20.1R2-2020-9-10.0_RELEASE_20.1R2_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: 20.1R2-2020-9-10.0_RELEASE_20.1R2_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.

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 have been developed at Juniper Networks over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information 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

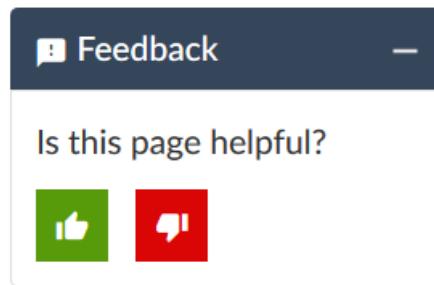
 **NOTE:** To obtain information about the components that are supported on the devices, and the special compatibility guidelines with the release, see the Hardware Guide for the product.

- **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

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.