

Release Notes

Published
2026-01-09

Junos OS Release 22.2R1

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX. These release notes accompany Junos OS Release 22.2R1. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can find release notes for all Junos OS releases at https://www.juniper.net/documentation/product/us/en/junos-os#cat=release_notes.

Table of Contents

Key Features in Junos OS Release 22.2 | 1

Junos OS Release Notes for ACX Series

What's New | 4

Class of Service | 4

Junos Telemetry Interface | 4

Routing Protocols | 5

Source Packet Routing in Networking (SPRING) or Segment Routing | 5

Additional Features | 6

What's Changed | 6

Known Limitations | 8

Open Issues | 8

Resolved Issues | 10

Migration, Upgrade, and Downgrade Instructions | 12

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 13

Junos OS Release Notes for cRPD

What's New | 14

MPLS | 14

What's Changed | 14

Known Limitations | 15

Open Issues | 15

Resolved Issues | 15

Junos OS Release Notes for cSRX

What's New | 16

Authentication and Access Control | 16

Network Address Translation (NAT) | 17

What's Changed | 17

Known Limitations | 17

Open Issues | 17

Resolved Issues | 18

Junos OS Release Notes for EX Series

What's New | 18

Hardware | 19

EVPN | 29

Licensing | 29

Routing Policy and Firewall Filters | 29

Additional Features | 30

What's Changed | 31

Known Limitations | 34

Open Issues | 35

Resolved Issues | 38

Migration, Upgrade, and Downgrade Instructions | 42

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 42

Junos OS Release Notes for JRR Series

What's New | 44

What's Changed | 44

Known Limitations | 44

Open Issues | 44

Resolved Issues | 44

Migration, Upgrade, and Downgrade Instructions | 45**Junos OS Release Notes for Juniper Secure Connect****What's New | 47****What's Changed | 47****Known Limitations | 47****Open Issues | 47****Resolved Issues | 47****Junos OS Release Notes for Junos Fusion for Enterprise****What's New | 48****What's Changed | 48****Known Limitations | 49****Open Issues | 49****Resolved Issues | 49****Migration, Upgrade, and Downgrade Instructions | 50****Junos OS Release Notes for Junos Fusion for Provider Edge****What's New | 56****What's Changed | 56****Known Limitations | 57****Open Issues | 57****Resolved Issues | 57****Migration, Upgrade, and Downgrade Instructions | 57****Junos OS Release Notes for MX Series****What's New | 67****What's New in 22.2R1-S2 | 67****Hardware | 68****Software Installation and Upgrade | 78**

User Authentication | 78

What's New in 22.2R1 | 79

EVPN | 80

High Availability | 81

Interfaces | 81

IP Tunneling | 82

Junos Telemetry Interface | 82

Licensing | 83

MACsec | 84

MPLS | 84

Platform and Infrastructure | 85

Precision Time Protocol (PTP) | 85

Routing Policy and Firewall Filters | 86

Routing Protocols | 86

Source Packet Routing in Networking (SPRING) or Segment Routing | 89

Software Installation and Upgrade | 89

Subscriber Management and Services | 90

VPNs | 91

Additional Features | 91

What's Changed | 93

Known Limitations | 100

Open Issues | 102

Resolved Issues | 109

Migration, Upgrade, and Downgrade Instructions | 125

Junos OS Release Notes for NFX Series

What's New | 131

Class of Service | 131

Unified Threat Management (UTM) | 131

Virtualized Network Functions (VNFs) | 132

What's Changed | 132

Known Limitations | 132

Open Issues | 133

Resolved Issues | 134

Migration, Upgrade, and Downgrade Instructions | 135

Junos OS Release Notes for PTX Series

What's New | 138

Interfaces | 138

IP Tunneling | 138

Junos Telemetry Interface | 138

Source Packet Routing in Networking (SPRING) or Segment Routing | 139

Routing Protocols | 140

Routing Policy and Firewall Filters | 141

Additional Features | 142

What's Changed | 142

Known Limitations | 148

Open Issues | 149

Resolved Issues | 151

Migration, Upgrade, and Downgrade Instructions | 153

Junos OS Release Notes for QFX Series

What's New | 158

Class of Service | 159

EVPN | 159

IP Tunneling | 161

MACsec | 161

Network Management and Monitoring | 162

Routing Protocols | 162

Routing Policy and Firewall Filters | 162

Additional Features | 163

What's Changed | 164

Known Limitations | 170

Open Issues | 171

Resolved Issues | 175

Migration, Upgrade, and Downgrade Instructions | 181

Junos OS Release Notes for SRX Series

What's New | 196

Class of Service | 196

Flow-Based and Packet-Based Processing | 196

High Availability | 197

J-Web | 197

Juniper Advanced Threat Prevention Cloud (ATP Cloud) | 198

MPLS | 198

Multicast | 199

Network Address Translation (NAT) | 199

Unified Threat Management (UTM) | 199

VPNs | 199

Additional Features | 200

What's Changed | 201

Known Limitations | 204

Open Issues | 204

Resolved Issues | 206

Migration, Upgrade, and Downgrade Instructions | 211

Junos OS Release Notes for vMX

What's New | 213

EVPN | 213

Junos Telemetry Interface | 213

MPLS | 213

OpenConfig | 214

Routing Protocols | 214

What's Changed | 215

Known Limitations | 217

Open Issues | 218

Resolved Issues | 218

Upgrade Instructions | 218

Junos OS Release Notes for vRR

What's New | 219

Routing Protocols | 219

What's Changed | 219

Known Limitations | 220

Open Issues | 220

Resolved Issues | 220

Junos OS Release Notes for vSRX

What's New | 221

Flow-Based and Packet-Based Processing | 222

Network Address Translation (NAT) | 222

Unified Threat Management (UTM) | 222

VPNs | 222

What's Changed | 223

Known Limitations | 225

Open Issues | 225

Resolved Issues | 226

Migration, Upgrade, and Downgrade Instructions | 228

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life
Releases | 234

Licensing | 235

Finding More Information | 236

Requesting Technical Support | 236

Revision History | 238

Key Features in Junos OS Release 22.2

Start here to learn about the key features in Junos OS Release 22.2. For more information about a feature, click the link in the feature description.

- **Support for dynamic address groups (cSRX)**—Starting in Junos OS Release 22.2R1, cSRX supports dynamic address groups (DAGs) or entries in a security policy.

In a Juniper Connected Security deployment, cSRX receives policy updates from external sources such as Policy Enforcer and SeclIntel feeds. These external sources provide lists of IP addresses that satisfy either of these conditions:

- Have a specific purpose, such as a blocklist.
- Include a common attribute, such as a particular location or behavior that might pose a threat.

You use the external intelligence in the cloud to identify threat sources by their IP addresses. You can then group those addresses into a dynamic address entry or DAG.

Reference this dynamic address entry in a security policy to control the traffic to and from those addresses.

[See [Dynamic Address Group Overview](#) and [Dynamic Address Groups in Security Policies](#).]

- **Automatically derived ESI configuration (MX Series, QFX5100, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX10002, QFX10002-60C, QFX10008, and QFX10016)**—In the current implementation, Junos OS derives the Ethernet segment identifier (ESI) from the system ID and the administrative key on the local multihomed provider edge (PE) device that is a part of the LACP link (actor). Starting in Junos OS Release 22.2R1, you can also configure the multihomed devices on an EVPN-VXLAN network to automatically generate the ESI from:

- The system ID and administrative key on the remote customer edge (CE) device (partner).
- The locally configured `mac` and `local-discriminator` values.

To automatically derive the ESI using the system ID and administrative key on the remote CE device, include `type-1-lacp` at the `[edit interfaces aeX aggregated-ether-options lacp auto-derive]` hierarchy level.

To automatically derive the ESI using locally configured values, configure `mac` and `local-discriminator` at the `[edit interfaces aeX aggregated-ether-options lacp auto-derive type-3-system-mac]` hierarchy level.

[See [Understanding Automatically Generated ESIs in EVPN Networks](#).]

- **Certificate-based authentication and encryption for MACsec (MX Series)**—Starting in Junos OS Release 22.2R1, you can enable MACsec on links connecting switches or routers using certificate-based authentication and encryption. Connected devices can mutually authenticate using 802.1X

over Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) and dynamically derive the connectivity association key (CAK) for encryption.

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

- **EVPN active/active redundancy, aliasing, and mass MAC withdrawal (MX Series and vMX)**—Starting in Junos OS Release 22.2R1, the listed devices support EVPN active/active redundancy, aliasing, and mass MAC withdrawal, integrated with VXLAN in the data plane. These features provide resilient inter-data center connectivity to the established Data Center Interconnect (DCI) technologies. This new support builds an end-to-end DCI solution by integrating EVPN active/active multicast with DP VXLAN.

Use existing configuration statements to configure active/active redundancy at the ESI level on the loopback (lo0) interface. Include lo0 as the virtual tunnel endpoint (VTEP) interface in the routing instance.

[See [EVPN-over-VXLAN Supported Functionality](#).]

- **NP-cache scale-up (SRX4600)**—Starting in Junos OS Release 22.2R1, the NP-cache wing count is 20 million. With this increment, the number of Express Path sessions increase fourfold.

[See [Sessions per Wing Statistics](#).]

- **Optimized intersubnet multicast (OISM) with MAC-VRF instances and IGMPv2 or IGMPv3 in an EVPN-VXLAN fabric (EX4650, QFX5110, QFX5120, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.2R1, you can configure OISM on leaf devices and border leaf devices in an EVPN-VXLAN ERB overlay fabric with:

- MAC-VRF routing instances or the default switch instance with IGMPv2 or IGMPv3.
- IGMP snooping and selective multicast Ethernet tag (SMET) forwarding optimizations with IGMPv2 or IGMPv3.

When you configure OISM, you must enable OISM and IGMP snooping on all the server leaf and border leaf devices in the EVPN-VXLAN fabric. With a MAC-VRF instance configuration, you configure the OISM supplemental bridge domain (SBD) and all revenue VLANs in the MAC-VRF instances on all leaf and border leaf devices in the fabric.

[See [Optimized Intersubnet Multicast in EVPN Networks](#).]

- **Support for guaranteed bit rate (GBR) on Junos Multi-Access User Plane (MX240, MX480, and MX960)**—Starting in Junos OS Release 22.2R1, the Junos Multi-Access User Plane has added GBR support and supports 3GPP standards for both 4G and 5G networks. The following features are added:
 - GBR support in the downlink direction and partial support in the uplink direction
 - Bandwidth reservation for express and GBR traffic flows

- Mapping of transport level marking to forwarding classes
- Call admission control (CAC)
- Maximum bit rate (MBR) and GBR policers

[See [QoS in Junos Multi-Access User Plane](#).]

- **Support for IPv6 tunnel (SRX Series and vSRX 3.0)**— Starting in Junos OS Release 22.2R1, you can encapsulate IPv4 and IPv6 traffic over the IPv6 network.

The IPv6 tunnel helps IPv4 traffic traverse over the IPv6 network. You can use IPv6 tunneling in various features such as policy routing and preferential billing. For example, a set-top box that supports only IPv4 traffic can traverse the server over an IPv6 network.

[See [show security flow session](#).]

- **Symmetric integrated routing and bridging (IRB) with EVPN Type 2 routes** (EX4400, EX4650, EX9204, EX9208, EX9214, MX Series, vMX, QFX5110, QFX5120, QFX10002, QFX10002-60C, QFX10008, and QFX10016). We support this feature only with MAC-VRF EVPN routing instance configurations and MAC-VRF service types `vlan-based` and `vlan-aware`. [See [Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics](#) and [irb-symmetric-routing](#).]

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [What's New | 4](#)
- [What's Changed | 6](#)
- [Known Limitations | 8](#)
- [Open Issues | 8](#)
- [Resolved Issues | 10](#)
- [Migration, Upgrade, and Downgrade Instructions | 12](#)

These release notes accompany Junos OS Release 22.2R1 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

- [Class of Service | 4](#)
- [Junos Telemetry Interface | 4](#)
- [Routing Protocols | 5](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 5](#)
- [Additional Features | 6](#)

Learn about new features introduced in this release for ACX Series routers.

Class of Service

- **Support for hierarchical scheduling on logical interface set (ACX5448, ACX5448-M, and ACX5448-D)**—Starting in Junos OS Release 22.2R1, the ACX5448 line of routers support hierarchical scheduling on a set of logical interfaces (IFL-SET). Applying hierarchical quality of service (HCoS) on a logical interface set enables four levels of scheduling on all the logical interfaces in that set. HCoS applies the levels one by one on all members of the set. You can configure TCP and TCP-REMAINING on the logical interface set as part of the HCoS configuration. You can also configure different CoS features through the CLI, as per your need and deployment model.

To enable hierarchical scheduling, set `hierarchical-scheduler` at the `[edit interfaces interface-name hierarchy-level]`.

See [[Hierarchical Class of Service in ACX Series Routers](#).]

Junos Telemetry Interface

- **Support for CPU state sensor (ACX710, ACX5448, MX204, MX240, MX150, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, and PTX10002)**—Starting in Junos OS Release 22.2R1, use the resource path `/system/cpus/cpu/state/` to export CPU parameters and including CPU usage per process and CPU usage per Routing Engine core information from a device to a collector.

[See [Telemetry Sensor Explorer](#).]

- **Network instance support enhancements (ACX710, ACX5448, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, and PTX10002)—** Starting in Junos OS Release 22.2R1, JTI supports new sensors for network instance statistics for the OpenConfig modules `openconfig-network-instance.yang` and `openconfig-routing-policy.yang`. The support includes OpenConfig configuration and streaming of state data.

[See [Telemetry Sensor Explorer](#) for telemetry support and [OpenConfig User Guide](#) for configuration.]

Routing Protocols

- **Nonstop active routing (NSR) support with BGP RIP sharding and BGP UpdateIO features (ACX5048, ACX5096, ACX5448, MX240, MX960, MX2008, MX10016, and PTX5000)—** Starting in Junos OS Release 22.2R1, we've enabled nonstop routing (NSR) for BGP RIP sharding and BGP UpdateIO features. With NSR enabled, the backup Routing Engine and backup routing protocol process (rpd) become the primary Routing Engine without negatively affecting the BGP peering sessions with the neighbors if the primary Routing Engine fails. The backup rpd processes the replicated BGP control-plane information and populates the route state in the same multithreaded manner as in the primary rpd.

After you configure NSR, the `show bgp neighbor` and `show bgp summary` commands display the information about the specific shards in the backup Routing Engine. To display the replicated information for a specific shard in the `show bgp replication` command, use the `rib-sharding shard-name` option.

See [[show bgp neighbor](#), [show bgp summary](#), [show bgp replication](#), and [BGP Overview](#).]

- **Support for BGP flow specification (ACX5448, ACX5448-M, and ACX5448-D)—** Starting in Junos OS Release 22.2R1, ACX5448 routers support BGP flow specification (BGP flowspec) filters based on the match conditions and actions. BGP flow specification filters support ingress IPv4 and IPv6 addresses. BGP flow specification filters internally creates implicit forwarding table filters or FTFs to mitigate DDoS attacks quickly. The BGP flow specification filters created in hardware have more precedence compared to the interface family `inet` and `inet6` filters (IFF and FTFs).

[See [Forwarding Traffic Using BGP Flow Specification DSCP Action](#).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Support for application-specific link attribute in OSPFv2 for segment routing traffic engineering (ACX753, ACX710, MX204, MX960, MX10008, and MX2020)—** Starting in Junos OS Release 22.2R1, you can advertise different te-attributes such as `te-metric`, `delay-metric`, or `admin-groups` for RSVP and flexible algorithms on the same link. This is done using flexible algorithm specific application-specific link attribute as defined in RFC 8920.

To configure flexible algorithm application-specific te-attribute, include the application-specific statement at the `[edit protocols ospf area interface]` hierarchy level and the `strict-asla-based-flex-algorithm` statement at the `[edit protocols ospf source-packet-routing]` hierarchy level.

[See [Understanding OSPF Flexible Algorithm for Segment Routing](#).]

- **BGP classful transport (CT) support for IPv6 and Segment Routing Traffic-Engineered (SR-TE) color-only support (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 22.2R1, we support BGP-CT with IPv6 and BGP service-routes with a color-only mapping community. We have also enhanced the transport-class configuration statement to provide strict resolution without falling back on best-effort tunnels.

[See [use-transport-class](#), [BGP Classful Transport \(BGP-CT\) with Underlying Colored SR-TE Tunnels Overview](#).]

Additional Features

Support for the following features has been extended to these platforms.

- **Supported transceivers, optical interfaces, and DAC cables** Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 6](#)
- [General Routing | 7](#)
- [Network Management and Monitoring | 7](#)
- [Routing Protocols | 7](#)
- [VPNs | 8](#)

Learn about what changed in this release for ACX Series.

Authentication and Access Control

- **SHA-1 password format deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX and vSRX)**—We've removed the sha1 option at the [edit system login password format] hierarchy level because SHA-1 is no longer supported for plain-text password encryption.

General Routing

- OpenConfig container names for Point-to-Multipoint per interface ingress and egress sensors are modified for consistency from signalling to signaling.

Network Management and Monitoring

- **Changes to the NETCONF <edit-config> RPC response (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the <edit-config> operation returns an error, the NETCONF server does not emit a <load-error-count> element in the RPC response. In earlier releases, the <edit-config> RPC response includes the <load-error-count> element when the operation fails.
- **DES deprecation for SNMPv3**—The Data Encryption Standard (DES) privacy protocol for SNMPv3 is deprecated due to weak security and vulnerability to cryptographic attacks. For enhanced security, configure the triple Data Encryption Standard (3DES) or the Advanced Encryption Standard (CFB128-AES-128 Privacy Protocol) as the encryption algorithm for SNMPv3 users.

[See [privacy-3des](#) and [privacy-aes128](#).]

- **Change in unnumbered-address support for GRE tunnel**—Starting in Junos OS Release 24.4R1, there is a behavioural change in unnumbered-address support for GRE tunnel with IPV6 family and display donor interface for both IPV4 and IPV6 families of GRE tunnel. You can view interface donor details under show interfaces hierarchy level.

[See [show interfaces](#).]

Routing Protocols

- **SSH TCP forwarding disabled by default**—We've disabled the SSH TCP forwarding feature by default to enhance security. To enable the SSH TCP forwarding feature, you can configure the allow-tcp-forwarding statement at the [edit system services ssh] hierarchy level.

In addition, we've deprecated the tcp-forwarding and no-tcp-forwarding statements at the [edit system services ssh] hierarchy level.

[See [services \(System Services\)](#).]

VPNs

- **Changes to show mvpn c-multicast and show mvpn instance outputs**—The FwdNh output field displays the multicast tunnel (mt) interface in the case of Protocol Independent Multicast (PIM) tunnels.

[See [show mvpn c-multicast](#).]

Known Limitations

Learn about known limitations in Junos OS Release 22.2R1 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

- On the dynamic addition of IFLs to ifset, non-CIR IFLs might get starved. [PR1656876](#)

Open Issues

IN THIS SECTION

- [General Routing | 9](#)

Learn about open issues in Junos OS Release 22.2R1 for ACX Series routers.

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

General Routing

- In ACX5000 devices, some next-hop routes do not get installed properly and generates the following error message in the LPM mode:

```
Failed to h/w update ip uc route entry
```

For the default route, if the route changes from ecmp to non-ecmp HOLD nexthop, the Packet Forwarding Engine gets into a corrupted ecmp nexthop.

[PR1365034](#)

- On ACX5448 devices, MAC learning or aging might stop in the Layer 2 domain after excessive MAC movements or continuous interface flaps. There might be unexpected flooding traffic when the issue occurs. [PR1480235](#)
- Due to BRCM KBP issue route, lookup might fail. [PR1533513](#)
- On ACX devices, traffic issue might be observed with downstream devices when you configure the Precision Time Protocol(PTP) (G.8275.1 PTP profile) along with PHY timestamping and Multiprotocol Label Switching (MPLS) terminated on 10G interface. The transit PTP ipv4 packets gets updated with incorrect Correction Factor(CF). The issue might be restored by disabling the PHY stamping. However, disabling might impact the PTP performance. [PR1612429](#)
- For ACX5448 VM Host-based platforms, starting with Junos OS 21.4R1 release or later, the ssh and root login are required for copying line card image from Junos VM to Linux host during installation. The ssh and root login are required during installation. Use the deny-password option instead of deny option as default root-login option under the ssh configuration to allow internal trusted communication. [PR1629943](#)
- On ACX5048 and ACX5096 devices, Junos OS does not support interface speed 10m on 1G interface. [PR1633226](#)
- Convergence time might be more than 60 seconds for IS-IS TILFA node protection testing. [PR1634033](#)
- In case of routing instance type EVPN or EVPN-VPWS, the system automatically creates one default routing instance apart from EVPN and/or EVPN-VPWS routing instance. In the output of the show snmp mib walk jnxVpnInfo command, the number of configured routing instances are always one more than the number of EVPN and/or EVPN-VPWS instances configured in the system.

[PR1659466](#)

- On ACX5448 and ACX710 devices, all types of delegated BFD sessions configured on routing-instance other than the default routing-instance might not come up. [PR1633395](#)

- On ACX5448 device, if a firewall has a log action and applied on the physical interface or lo0 interface, the LDP neighbor cannot go up. [PR1648968](#)
- In case of routing instance type EVPN or EVPN-VPWS, system automatically creates one default routing instance apart from EVPN and/or EVPN-VPWS routing instance. Hence, in the output of the show snmp mib walk jnxVpnInfo command, the number of configured routing instances are always one more than number of EVPN and/or EVPN-VPWS instances configured in the system. [PR1659466](#)
- Some of the interfaces get zero status in the output of the monitor interface traffic command. Sending traffic across all interfaces and applying speed of 100m on all 1g copper ports, clears the interfaces status. [PR1661617](#)

Resolved Issues

IN THIS SECTION

- General Routing | [10](#)
- Platform and Infrastructure | [12](#)
- Routing Protocols | [12](#)

Learn about the issues fixed in this release for ACX Series.

General Routing

- On ACX5448 devices, the cfmd process might generate core file if you change the CCM configuration from the aggregated Ethernet interface ifl to physical ifl, and if the physical ifl was previously part of the aggregated Ethernet interface bundle. [PR1612212](#)
- On ACX5000 device, the Local fault and Remote fault signaling do not get logged on the `/var/log/messages` file. [PR1624761](#)
- Unicast might lose packet due to control-word configuration. [PR1626058](#)
- On ACX2000 devices, the output packet statistics does not get incremented on the unit even after configuring statistics. [PR1627040](#)

- Multicast traffic might drop if you enable the IGMP snooping for VLAN. [PR1628600](#)
- On ACX710 devices, when the panic command gets issued, box gets stuck and no vmcore file gets formed. [PR1629700](#)
- Late drops do not get at par with the PN configured. [PR1630724](#)
- The storm-control rate-limit might not work with VPLS policer under IFL. [PR1633427](#)
- DHCP clients might not come online for the IRB+VLAN/EVPN scenario. [PR1633778](#)
- IS-IS last transition time never increments. [PR1634747](#)
- On ACX5448 and ACX710 devices, the IPv6 BFD session over the aggregated Ethernet interface might remain down. [PR1635020](#)
- On ACX710 and ACX5448 devices working as a PE device stops forwarding the Layer 3 VPN traffic after core-facing link flaps. [PR1635801](#)
- On ACX5448 devices, the locally switched traffic might be dropped with ESI configured. [PR1638386](#)
- On ACX5448 and ACX710 devices, the Layer 3 interface creation might fail. [PR1638581](#)
- On ACX710 devices, the following message does not get generated when you use the USB image to recover the box:

USB Installation is done, Please remove the USB media to continue

[PR1640143](#)

- On ACX5448 devices, high priority packets might be dropped. [PR1642187](#)
- On ACX5448 devices, reboot reason that gets displayed is not as expected. [PR1643781](#)
- Due to the MAC learning limit being exceeded, traffic might get dropped in the MC-AE scenario. [PR1653926](#)
- The ARP request packets might be sent out from the ACX router without a VLAN header. [PR1638421](#)
- KRT queue entries gets stuck during the Routing Engine switchover when the backup RPD is not ready. [PR1641297](#)
- The LDP sessions might flap in the VPLS scenario resulting in the Packet Forwarding Engine errors. [PR1654172](#)
- The copper ports on ACX5448 device might go down if loaded with copper SFP. [PR1643989](#)

- Traffic might get silently discarded in the MPLS scenario with explicit-null. [PR1646097](#)
- While sending BGP notification messages for the RFC 8538 hard reset, the data portion sometimes are not present. [PR1648479](#)

Platform and Infrastructure

- The vmxt_lnx process generates core file at topo_get_link jnh_features_get_jnh jnh_stream_attach. [PR1638166](#)

Routing Protocols

- IPv6 inline BFD sessions are down when neighbor does not get resolved. [PR1650677](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | [13](#)

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

Starting in Junos OS release 21.2R1, all Junos OS products which were previously running on FreeBSD 11.x based Junos OS are migrated to FreeBSD 12.x based Junos OS, except EX4400. Starting with Junos OS release 21.3R1, EX4400 platforms are migrated to FreeBSD 12.x based Junos OS.

For information about software installation and upgrade, see the https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/software-installation-and-upgrade/software-installation-and-upgrade.html Installation and Upgrade Guide.

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

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

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

Table 1: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for cRPD

IN THIS SECTION

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

These release notes accompany Junos OS Release 22.2R1 for the containerized routing protocol process (cRPD) container. 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

- [MPLS | 14](#)

Learn about new features introduced in this release for cRPD.

MPLS

- **Support for RSVP delay constraint (cRPD, MX960, and PTX10008)**—Starting in Junos OS Release 22.2R1, you can configure RSVP label-switched paths (LSPs) to use a delay metric for computing the path. To configure, use the new CLI options that we've introduced under the [edit protocols mpls label-switched-path *name*] hierarchy. We've also updated the outputs of the following show commands:
 - show ted link detail
 - show ted database extensive
 - show route protocol bgp table lsdist.0 extensive
 - show spring-traffic-engineering lsp detail
 - show express-segments name *name* detail
 - show mpls lsp detail

What's Changed

There are no changes in behavior and syntax in Junos OS Release 22.2R1 for cRPD.

Known Limitations

There are no known limitations in hardware and software in Junos OS 22.2R1 for cRPD.

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

Open Issues

There are no known issues in hardware and software in Junos OS Release 22.2R1 for cRPD.

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

Resolved Issues

Learn about the issues fixed in this release for cRPD.

Routing Protocols

- Passive BGP session in no-forwarding instance could not come up. [PR1645010](#)

User Interface and Configuration

- Unable to access configure exclusive mode after mgd process is stopped. [PR1641025](#)

Junos OS Release Notes for cSRX

IN THIS SECTION

- [What's New | 16](#)
- [What's Changed | 17](#)
- [Known Limitations | 17](#)

- [Open Issues | 17](#)
- [Resolved Issues | 18](#)

These release notes accompany Junos OS Release 22.2R1 for the cSRX Container Firewall, a containerized version of the SRX Series Services Gateway. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [Authentication and Access Control | 16](#)
- [Network Address Translation \(NAT\) | 17](#)

Learn about new features introduced in this release cSRX.

Authentication and Access Control

- **Support for dynamic address groups (cSRX)**—Starting in Junos OS Release 22.2R1, cSRX supports dynamic address groups (DAGs) or entries in a security policy.

In a Juniper Connected Security deployment, cSRX receives policy updates from external sources such as Policy Enforcer and Seclntel feeds. These external sources provide lists of IP addresses that satisfy either of these conditions:

- Have a specific purpose, such as a blocklist.
- Include a common attribute, such as a particular location or behavior that might pose a threat.

You use the external intelligence in the cloud to identify threat sources by their IP addresses. You can then group those addresses into a dynamic address entry or DAG.

Reference this dynamic address entry in a security policy to control the traffic to and from those addresses.

[See [Dynamic Address Group Overview](#) and [Dynamic Address Groups in Security Policies](#).]

Network Address Translation (NAT)

- **NAT support for DNS (SRX Series, vSRX, and cSRX)**—Starting in Junos OS Release 22.2R1, you can use DNS and a fully qualified domain name (FQDN) with either source NAT or destination NAT as part of your NAT configuration.

You can use DNS name servers to resolve hostnames to IP addresses. A DNS cache time to live (TTL) is introduced under the address-book option for each DNS name entry. We support a minimum DNC cache TTL of 16 seconds.

In case of multiple IP addresses in the DNS response, the first IP address in the response is added to the NAT pool.

[See [Address Books and Address Sets](#) and [show security nat source pool](#).]

What's Changed

There are no changes in behavior and syntax in Junos OS Release 22.2R1 for cSRX.

Known Limitations

There are no known limitations in hardware and software in Junos OS 22.2R1 for cSRX.

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

Open Issues

There are no known issues in hardware and software in Junos OS Release 22.2R1 for cSRX.

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

Resolved Issues

There are no resolved issues in Junos OS Release 22.2R1 for cSRX.

Junos OS Release Notes for EX Series

IN THIS SECTION

- [What's New | 18](#)
- [What's Changed | 31](#)
- [Known Limitations | 34](#)
- [Open Issues | 35](#)
- [Resolved Issues | 38](#)
- [Migration, Upgrade, and Downgrade Instructions | 42](#)

These release notes accompany Junos OS Release 22.2R1 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

- [Hardware | 19](#)
- [EVPN | 29](#)
- [Licensing | 29](#)
- [Routing Policy and Firewall Filters | 29](#)

● Additional Features | 30

Learn about new features introduced in this release for EX Series switches.

Hardware

- **New EX4100 and EX4100-F switches**—Starting in Junos OS Release 22.2R1, we introduce the EX4100 and EX4100-F family of switches that provide connectivity for high-density environments and scalability for network growth. You can deploy the EX4100 and EX4100-F stackable switches in small, medium, and large campus and branch enterprise networks. We support 24-port and 48-port switch variants with or without PoE+ and with different airflow directions. The switches have dedicated Virtual Chassis ports (VCPs) and uplink ports.

We support the following switches: EX4100-48P, EX4100-48T, EX4100-48T-AFI, EX4100-48T-DC, EX4100-24P, EX4100-24T, EX4100-24T-DC, EX4100-F-48P, EX4100-F-48T, EX4100-F-24P, and EX4100-F-24T.

Table 2: Features Supported on EX4100 and EX4100-F Switches

Feature	Description
Access and authentication	<ul style="list-style-type: none">• FQDN support in RADIUS configuration. The RADIUS server configuration supports fully qualified domain names (FQDN) that resolve to one or more IP addresses. [See Specifying RADIUS Server Connections on Switches.]• 802.1X authentication. [See 802.1X Authentication.] <p>Captive portal. [See Captive Portal Authentication.]</p>

Table 2: Features Supported on EX4100 and EX4100-F Switches (*Continued*)

Feature	Description
Chassis	<ul style="list-style-type: none"> FRU management and environment monitoring, and chassis support for EX4100 switches only, including: <ul style="list-style-type: none"> PSU, fan, and temperature sensors monitoring Power management support for two power supply units (PSUs) and two field-replaceable fans. The system functions with one fan until it reaches shutdown temperature. When temperature reported by various sensors crosses the specified threshold, the fan speed increases or decreases to regulate the temperature. If the temperature exceeds the shutdown threshold, system shutdown is initiated. <p>[See Understanding Power Management on EX Series Switches.]</p>
CoS	<ul style="list-style-type: none"> Support for CoS configuration. <p>[See Junos OS CoS for EX Series Switches Overview.]</p>

Table 2: Features Supported on EX4100 and EX4100-F Switches (*Continued*)

Feature	Description
EVPN	<ul style="list-style-type: none"> Support for EVPN-VXLAN group-based policies. EX4100 and EX4100-F switches provide standards-based multilevel segmentation (also called group-based policy, or GBP) on the basis of Layer 3 virtual networks and group-based tags rather than IP-based filters. This support allows for different levels of access control for endpoints and applications even within the same VLAN. The EX4100 and EX4100-F switches also provide GBP support for locally switched traffic on VXLAN access ports. <p>[See Micro and Macro Segmentation using Group Based Policy in a VXLAN.]</p> <ul style="list-style-type: none"> Support for the following Layer 2 VXLAN gateway services in an EVPN-VXLAN network: <ul style="list-style-type: none"> 802.1X authentication, accounting, central web authentication (CWA) authentication, and captive portal CoS DHCPv4 and DHCPv6 snooping, dynamic ARP inspection (DAI), neighbor discovery inspection, IP source guard and IPv6 source guard, and router advertisement (RA) guard (no multihoming) Firewall filters and policing Storm control, port mirroring, and MAC filtering <p>[See EVPN Feature Guide.]</p> <ul style="list-style-type: none"> Support for Layer 3 VXLAN gateway in EVPN-VXLAN centrally routed bridging (CRB) overlay or edge-routed bridging (ERB) overlay networks on standalone switches or Virtual Chassis. The switch supports the following features: <ul style="list-style-type: none"> Default gateway using IRB interfaces to route traffic between VLANs. [See Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network.

Table 2: Features Supported on EX4100 and EX4100-F Switches (Continued)

Feature	Description
	<ul style="list-style-type: none"> • IPv6 data traffic routed through an EVPN-VXLAN overlay network with an IPv4 underlay. [See Routing IPv6 Data Traffic through an EVPN-VXLAN Network with an IPv4 Underlay.] • EVPN pure Type 5 routes. [See Understanding EVPN Pure Type-5 Routes.] <p>The Virtual Chassis doesn't support EVPN-VXLAN multihoming, but you can use the standalone switch as an EVPN-VXLAN provider edge (PE) device in multihoming use cases. We support the following Layer 2 VXLAN gateway features in an EVPN-VXLAN network:</p> <ul style="list-style-type: none"> • Active/active multihoming • Proxy ARP use and ARP suppression, and Neighbor Discovery Protocol (NDP) use and NDP suppression on non-IRB interfaces • Ingress node replication for broadcast, unknown unicast, and multicast (BUM) traffic forwarding <p>[See EVPN Feature Guide.]</p>
Flow monitoring	<ul style="list-style-type: none"> • Support for flow-based telemetry –You can configure flow-based telemetry (FBT) and additional parameters to track for a flow using the feature-profile <i>name</i> features statement at the [edit inline-monitoring] hierarchy level. <p>See [features and Flow-Based Telemetry (EX4100, EX4100-F, and EX4400 Series).]</p>

Table 2: Features Supported on EX4100 and EX4100-F Switches (*Continued*)

Feature	Description
Hardware	<ul style="list-style-type: none"> • New EX4100 and EX4100-F switch models— We introduce the following models of the EX4100 Ethernet Switches: <ul style="list-style-type: none"> • EX4100-24P, EX4100-24T, and EX4100-24T-DC –Twenty-four 10/100/1000-Mbps RJ-45 ports, four 10/25-Gbps SFP28 Virtual Chassis ports (VCPs), and four 1000-Mbps/10-Gbps SFP+ uplink ports on the front panel. Only EX4100-24P has PoE+ enabled ports. EX4100-24T-DC is powered by DC power supplies; the rest of the switch models are powered by AC power supplies. All these switch models have AFO cooling. • EX4100-48P, EX4100-48T, EX4100-48T-AFI, EX4100-48T-DC—Forty-eight 10/100/1000-Mbps RJ-45 ports, four 10/25-Gbps SFP28 Virtual Chassis ports, and four 1000-Mbps/10Gbps SFP+ uplink ports on the front panel. Only EX4100-48P has PoE+ enabled ports. EX4100-48T-DC is powered by DC power supplies; the rest of the switch models are powered by AC power supplies. EX4100-48T-AFI has AFI cooling; the other switch models have AFO cooling. • EX4100-F-24P and EX4100-F-24T—Twenty-four 10/100/1000-Mbps RJ-45 ports, four 1/10 Gbps SFP+ Virtual Chassis ports, and four 1000-Mbps/10 Gbps SFP+ uplink ports on the front panel. Only EX4100-F-24P has PoE+ enabled ports. The switch models are powered by built-in AC power supplies and built-in AFO cooling. • EX4100-F-48P and EX4100-F-48T—Forty-eight 10/100/1000-Mbps RJ-45 ports, four 1/10 Gbps SFP+ Virtual Chassis ports, and four 1000-Mbps/10 Gbps SFP+ uplink ports on the front panel. Only EX4100-F-48P has PoE+ enabled ports. The switch models are powered by built-in AC power supplies and built-in AFO cooling.
High availability and resiliency	<ul style="list-style-type: none"> • Resiliency support for inter-integrated controller (I2C), disk failure, and disk health. <p>[See High Availability User Guide.]</p>

Table 2: Features Supported on EX4100 and EX4100-F Switches (*Continued*)

Feature	Description
Interfaces	<ul style="list-style-type: none"> One multi-rate FPC and three multi-rate PICs. <p>EX4100-48P, EX4100-48T, EX4100-24P, and EX4100-24T support the following speeds:</p> <ul style="list-style-type: none"> Downlink ports on PIC 0 (ports 0–47 on EX4100-48P and EX4100-48T, ports 0–23 on EX4100-24P and EX4100-24T) support 10-Mbps, 100-Mbps, and 1-Gbps speeds. VCPs (ports 0–3 on PIC 1) support 4x10-Gbps or 4x25-Gbps speeds. If you convert the VCPs to uplink ports, ports 0 through 3 on PIC1 support 1-Gbps speeds. Uplink ports (ports 0–3 on PIC 2) support 4x10-Gbps or 4x1-Gbps speeds. <p>EX4100-F-48P, EX4100-F-48T, EX4100-F-24P, and EX4100-F-24T support the following speeds:</p> <ul style="list-style-type: none"> Downlink ports on PIC 0 (ports 0–47 for EX4100-F-48P and EX4100-F-48T, ports 0–23 for EX4100-F-24P and EX4100-F-24T) support 10-Mbps, 100-Mbps, and 1-Gbps speeds. VCPs (ports 0–3 on PIC 1) support 4x10-Gbps speeds. If you convert the VCPs to uplink ports, ports 0 through 3 on PIC1 support 1-Gbps speeds. Uplink ports (ports 0–3 on PIC 2) support 4x10-Gbps or 4x1-Gbps speeds. <p>[See Port speed.]</p> <ul style="list-style-type: none"> Optics support. [See Hardware Compatibility Tool.] PoE support. EX4100 and EX4100-F switches support 802.3AT PoE+, fast PoE, and perpetual PoE . <p>[See Understanding PoE on EX Series Switches.]</p>

Table 2: Features Supported on EX4100 and EX4100-F Switches (Continued)

Feature	Description
Junos telemetry interface	<ul style="list-style-type: none"> Support for JTI Packet Forwarding Engine and Routing Engine sensor. You can use the Junos telemetry interface (JTI) and remote procedure calls (gRPC) to stream statistics from the switches to an outside collector. Support for secure packet capture to Cloud using JTI. You can use Junos telemetry interface (JTI) to capture packets from a device and send them over a secure channel to an external collector (in the cloud) for monitoring and analysis. <p>To use secure packet capture, include the /junos/system/linecard/packet-capture resource path using a Junos remote procedure call (RPC).</p>
Layer 2 features	<ul style="list-style-type: none"> Support for Layer 2 features. [See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation, Understanding Layer 2 Bridge Domains, and Understanding Layer 2 Learning and Forwarding.] Support for Layer 2 multicast features. [See Multicast Overview and Understanding Multicast Snooping. Use the interface-name and ip-address options to configure the management address on the switch. [See Configuring LLDP (CLI Procedure) .]
Layer 3 features	<ul style="list-style-type: none"> Support for Layer 3 features and interior gateway protocols (OSPF, IS-IS, RIP, and ECMP) for IPv4 and IPv6. [See Understanding OSPF Configurations and BGP Overview.]

Table 2: Features Supported on EX4100 and EX4100-F Switches (Continued)

Feature	Description
Licensing	<ul style="list-style-type: none"> You need a license to use the software features on the EX4100 and EX4100-F switches. To know more about licenses and supported features, see Flex Software License for EX Series Switches. <p>To add, delete, and manage licenses, see Managing Licenses.</p>
Network management and monitoring	<ul style="list-style-type: none"> Support for Ethernet Operation, Administration, and Maintenance (OAM) and VRRP. [See Ethernet OAM and CFM for Switches.] Support for IEEE 802.1ag CFM on service provider interfaces and Q-in-Q (point-to-point) interfaces. [See Introduction to OAM Connectivity Fault Management (CFM).] Support for Juniper Mist Wired Assurance. You can automatically onboard the EX4100 and EX4100-F switches to the Juniper Mist Cloud using a single activation code and provision the switch interfaces. [See Juniper AI-Driven Enterprise and Overview of EX Series Switches and the Juniper Mist Cloud.] Support for: <ul style="list-style-type: none"> Spanning-tree protocols. [See Spanning Tree Protocol Instances and Interfaces.] sFlow network monitoring technology. [See sFlow Monitoring Technology.] Local and remote port mirroring, and remote port mirroring to an IP address (GRE encapsulation). [See Port Mirroring and Analyzers.]

Table 2: Features Supported on EX4100 and EX4100-F Switches (*Continued*)

Feature	Description
Software installation and upgrade	<ul style="list-style-type: none"> Support for DHCP option 43 suboption 8 to provide proxy server information in phone-home client. During the bootstrapping process, the phone-home client (PHC) can access the redirect server through a proxy server. The DHCP server uses DHCP option 43 suboption 8 to deliver the details of IPv4 and/or IPv6 proxy servers to the PHC. The DHCP daemon running on the target switch learns about the proxy servers in the initial DHCP cycle and then populates either the <code>phc_vendor_specific_info.xml</code> or the <code>phc_v6_vendor-specific_info.xml</code> file located in the <code>/var/etc/</code> directory with the vendor-specific information. Support for the phone-home client. The phone-home client (PHC) can securely provision an EX4100 or EX4100-F Virtual Chassis without requiring user interaction. You only need to: <ul style="list-style-type: none"> Ensure that the Virtual Chassis members have the factory-default configuration. Interconnect the member switches using dedicated or default-configured VCPs. Connect the Virtual Chassis management port or any network port to the network. Power on the Virtual Chassis members. The PHC automatically starts up on the Virtual Chassis and connects to the phone-home server (PHS). The PHS responds with bootstrapping information, including the Virtual Chassis topology, software image, and configuration. The PHC upgrades each Virtual Chassis member with the new image and applies the configuration, and the Virtual Chassis is ready to go. <p>[See Obtaining Configurations and Software Image Without User Intervention Using Phone-Home Client.]</p> Secure boot support in U-boot phase to authenticate and verify the loaded software image while also preventing software-based attack.

Table 2: Features Supported on EX4100 and EX4100-F Switches (Continued)

Feature	Description
	<p>[See Software Installation and Upgrade Guide.]</p> <ul style="list-style-type: none"> • ZTP with IPv6. You can use either the legacy DHCP-options-based zero-touch provisioning (ZTP) or the phone-home client (PHC) to provision software for the EX4100 and EX4100-F switches. If the switch boots up and receives DHCP options from the DHCP server for ZTP, ZTP resumes. If DHCP options are not present, the switch attempts the PHC method. <p>The DHCP server uses DHCPv6 options 59 and 17 and applicable suboptions to exchange ZTP-related information between itself and the DHCP client.</p> <p>[See Zero Touch Provisioning Overview.]</p>
Timing	<ul style="list-style-type: none"> • Support for Precision Time Protocol (PTP) transparent clock on uplink ports connected to external MACsec PHY (EX4100-48 and EX4100-24). <p>[See Understanding Transparent Clocks in Precision Time Protocol.]</p> <ul style="list-style-type: none"> • Support for PTP transparent clock for all ports (EX4100-F-48 and EX4100-F-24) when MACsec is not enabled. <p>[See Understanding Transparent Clocks in Precision Time Protocol.]</p>
Uplink failure detection	<ul style="list-style-type: none"> • Support for debounce interval configuration. You can configure the debounce interval, which is the time (in seconds) that elapses before the downlink interfaces are brought up after a state change of the uplink interfaces. <p>You configure the debounce-interval statement at the [edit protocols uplink-failure-detection group <i>group-name</i>] hierarchy level.</p> <p>[See Uplink Failure Detection.]</p>

Table 2: Features Supported on EX4100 and EX4100-F Switches (Continued)

Feature	Description
Virtual Chassis	<ul style="list-style-type: none"> Support for Virtual Chassis configuration. You can interconnect an EX4100 or EX4100 Multigigabit or EX4100-F switch with other EX4100 or EX4100-F switches into a Virtual Chassis in non-mixed mode. <p>[See Virtual Chassis Overview for Switches.]</p>

EVPN

- Optimized intersubnet multicast (OISM) with MAC-VRF instances and IGMPv2 or IGMPv3 in an EVPN-VXLAN fabric (EX4650, QFX5110, QFX5120, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.2R1, you can configure OISM on leaf devices and border leaf devices in an EVPN-VXLAN ERB overlay fabric with:
 - MAC-VRF routing instances or the default switch instance with IGMPv2 or IGMPv3.
 - IGMP snooping and selective multicast Ethernet tag (SMET) forwarding optimizations with IGMPv2 or IGMPv3.

When you configure OISM, you must enable OISM and IGMP snooping on all the server leaf and border leaf devices in the EVPN-VXLAN fabric. With a MAC-VRF instance configuration, you configure the OISM supplemental bridge domain (SBD) and all revenue VLANs in the MAC-VRF instances on all leaf and border leaf devices in the fabric.

[See [Optimized Intersubnet Multicast in EVPN Networks](#).]

Licensing

- Juniper Agile Licensing (EX4100)**—Starting in Junos OS Release 22.2R1, the EX4100 switch support Juniper Agile Licensing.

Juniper Agile Licensing provides simplified and centralized license administration and deployment. You can use Juniper Agile Licensing to install and manage licenses for hardware and software features.

[See [Flex Software License for EX Series Switches](#) and [Managing Licenses](#).]

Routing Policy and Firewall Filters

- Optimize TCAM when EVPN/VXLAN is enabled (EX4400-48F, EX4650, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200, and QFX5210)**—

In Junos OS Release 22.2R1, we've introduced CLI configuration commands to optimize ternary content addressable memory (TCAM) space usage. Use these commands to prevent ingress filter processor (IFP) TCAM space exhaustion:

- `set chassis ivacl-firewall-no-portrange-profile`
- `set chassis iracl-firewall-ipv4-profile`
- `set chassis ipvacl-firewall-l2-profile`
- `set chassis input-firewall-optimized-profile`

Additional Features

Support for the following features has been extended to these platforms.

- **Lightweight PE-CE Loop Detection on EVPN-VXLAN Fabrics (EX4400-48MP, EX4400-48P, EX9200, MX240, MX480, MX960, and MX10003)**

[See [EVPN-VXLAN Lightweight Leaf to Server Loop Detection](#).]

- **Support for BPDU protection for EVPN-VXLAN (EX4300-48MP, EX4400, EX4650, QFX5110, QFX5120, QFX10002-60C, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.2R1, we support bridge protocol data unit (BPDU) protection for EVPN-VXLAN.

[See [Understanding BPDU Protection for EVPN-VXLAN](#).]

- **Support for EVPN routing policies (EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, and EX4400-48T)**—We've extended support of policy filter configurations for EVPN routes to the listed EX4400 switches.

[See [Routing policies for EVPN](#) .]

- **Supported transceivers, optical interfaces, and DAC cables** Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

- **Symmetric integrated routing and bridging (IRB) with EVPN Type 2 routes** (EX4400, EX4650, EX9204, EX9208, EX9214, MX Series, vMX, QFX5110, QFX5120, QFX10002, QFX10002-60C, QFX10008, and QFX10016). We support this feature only with MAC-VRF EVPN routing instance configurations and MAC-VRF service types `vlan-based` and `vlan-aware`. [See [Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics](#) and [irb-symmetric-routing](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 31](#)
- [General Routing | 31](#)
- [MPLS | 32](#)
- [Network Management and Monitoring | 32](#)
- [Routing Protocols | 33](#)
- [User Interface and Configuration | 33](#)
- [VPNs | 34](#)

Learn about what changed in this release for EX Series.

Authentication and Access Control

- **SHA-1 password format deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX and vSRX)**—We've removed the sha1 option at the [edit system login password format] hierarchy level because SHA-1 is no longer supported for plain-text password encryption.

General Routing

- **Instance type change is not permitted from default to L3VRF in open configuration (ACX Series, EX Series, MX Series, QFX Series, SRX Series, vMX, and vSRX)**—DEFAULT_INSTANCE is the primary instance that runs when there is no specific instance type configured in the route set routing-options. Any instance you explicitly configure is translated into set routing-instance r1 routing-options. The issue appears in translation, when you change instance type DEFAULT_INSTANCE (any instance to DEFAULT_INSTANCE) to L3VRF or L3VRF to DEFAULT_INSTANCE. As a result, such changes are not permitted. Additionally, DEFAULT_INSTANCE can only be named DEFAULT, and DEFAULT is reserved for DEFAULT_INSTANCE, therefore allowing no such changes.
- OpenConfig container names for Point-to-Multipoint per interface ingress and egress sensors are modified for consistency from "signalling" to "signaling".

MPLS

- Starting with Junos OS 16.1 the MPLS EXP bits transmitted in self ping messages are set based on the DSCP/ToS setting of the corresponding IP packet.
- When defining a constrained path LSP using more than one strict hop belonging to the egress node, the first strict hop must be set to match the IP address assigned to the egress node on the interface that receives the RSVP Path message. If the incoming RSVP Path message arrives on an interface with a different IP address the LSP is rejected.

Network Management and Monitoring

- **DES deprecation for SNMPv3**—The Data Encryption Standard (DES) privacy protocol for SNMPv3 is deprecated due to weak security and vulnerability to cryptographic attacks. For enhanced security, configure the triple Data Encryption Standard (3DES) or the Advanced Encryption Standard (CFB128-AES-128 Privacy Protocol) as the encryption algorithm for SNMPv3 users.

[See [privacy-3des](#) and [privacy-aes128](#).]

- **Changes when deactivating or deleting instances of the ephemeral configuration database (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The following changes apply when you deactivate or delete ephemeral database instances in the static configuration database:
 - When you deactivate the entire [edit system configuration-database ephemeral] hierarchy level, the device deletes the files and corresponding configuration data for all user-defined ephemeral instances. In earlier releases, the files and configuration data are preserved; however, the configuration data is not merged with the static configuration database.
 - When you delete an ephemeral instance in the static configuration database, the instance's configuration files are also deleted. In earlier releases, the configuration files are preserved.
 - You can delete the files and corresponding configuration data for the default ephemeral database instance by configuring the delete-ephemeral-default statement in conjunction with the ignore-ephemeral-default statement at the [edit system configuration-database ephemeral hierarchy level.

[See [Enable and Configure Instances of the Ephemeral Configuration Database](#).]

- **Changes to the NETCONF <edit-config> RPC response (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the <edit-config> operation returns an error, the NETCONF server does not emit a <load-error-count> element in the RPC response. In earlier releases, the <edit-config> RPC response includes the <load-error-count> element when the operation fails.

- **Change in unnumbered-address support for GRE tunnel**—Starting in Junos OS Release 24.4R1, there is a behavioural change in unnumbered-address support for GRE tunnel with IPV6 family and display donor interface for both IPV4 and IPV6 families of GRE tunnel. You can view interface donor details under show interfaces hierarchy level.

[See [show interfaces](#).]

- **Deprecated REST API Ciphers Unsupported in OpenSSL 3.0**—We have deprecated ciphers unsupported in the OpenSSL3.0 version. We also have introduced support for new ciphers introduced in TLSv1.3. for Rest API.

Routing Protocols

- The RPD OSPF LDP SYNC message not logged?On all Junos OS and Junos OS Evolved devices, when an LDP session goes down there is a loss of synchronization between LDP and OSPF. After the loss of synchronization, when an interface has been in the holddown state for more than three minutes, the system log message with a warning level is sent. This message appears in both the messages file and the trace file. However, the system log message does not get logged if you explicitly configure the `hold-time` for ldp-synchronization at the `edit protocols ospf area area id interface interface name` hierarchy level less than three minutes. The message is printed after three minutes.
- **SSH TCP forwarding disabled by default**—We've disabled the SSH TCP forwarding feature by default to enhance security. To enable the SSH TCP forwarding feature, you can configure the `allow-tcp-forwarding` statement at the `edit system services ssh` hierarchy level.In addition, we've deprecated the `tcp-forwarding` and `no-tcp-forwarding` statements at the `edit system services ssh` hierarchy level.

See [[services \(System Services\)](#).]

User Interface and Configuration

- **Load JSON configuration data with unordered list entries (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos schema requires that list keys precede any other siblings within a list entry and appear in the order specified by the schema. Junos devices provide two options to load JSON configuration data that contains unordered list entries:
 - Use the `request system convert-json-configuration` operational mode command to produce JSON configuration data with ordered list entries before loading the data on the device.
 - Configure the `reorder-list-keys` statement at the `[edit system configuration input format json]` hierarchy level. After you configure the statement, you can load JSON configuration data with

unordered list entries, and the device reorders the list keys as required by the Junos schema during the load operation.

- When you configure the `reorder-list-keys` statement, the load operation can take significantly longer to parse the configuration, depending on the size of the configuration and number of lists. Therefore, for large configurations or configurations with many lists, we recommend using the `request system convert-json-configuration` command instead of the `reorder-list-keys` statement.

[See [json](#) and [request system convert-json-configuration](#)]

- Junos XML protocol Perl modules deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—We no longer provide the Junos XML protocol Perl client for download. To use Perl to manage Junos devices, use the NETCONF Perl library instead.

[See [Understanding the NETCONF Perl Client and Sample Scripts..](#)]

- When you configure `max-cli-sessions` at the `[edit system]` hierarchy level, it restricts the maximum number of cli sessions that can coexist at any time. Once the `max-cli-sessions` number is reached, new CLI access is denied. The users who are configured to get the CLI upon login, are also denied new login. The `max-cli-sessions` is configured so you can control the memory usage for the CLI. You may set the `max-cli-sessions` per your requirement. However, if `max-cli-sessions` is not configured, there is no control on the number of CLIs getting invoked.

VPNs

- Changes to `show mvpn c-multicast` and `show mvpn instance outputs`**—The `FwdNh` output field displays the multicast tunnel (mt) interface in the case of Protocol Independent Multicast (PIM) tunnels.

[See [show mvpn c-multicast](#).]

Known Limitations

IN THIS SECTION

- Platform and Infrastructure | 35

Learn about known limitations in Junos OS Release 22.2R1 for EX Series switches.

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

Platform and Infrastructure

- In EVPN_VXLAN deployment, BUM traffic replication over VTEP might send out more packets than expected. [PR1570689](#)
- On EX4300-MP devices, when the command request system software rollback is performed device is going down and dcpfe core files are generated. [PR1631640](#)
- Redirect server doesn't handle cluster ID requests. [PR1646141](#)
- POE is supported only on down link ports. If the user configures the same POE on other ports which are not supported, then configuration will go through. However, POE won't work on those ports. [PR1647143](#)

Open Issues

IN THIS SECTION

- [Forwarding and Sampling | 36](#)
- [Infrastructure | 36](#)
- [Layer 2 Features | 36](#)
- [Network Management and Monitoring | 36](#)
- [Platform and Infrastructure | 36](#)
- [Virtual Chassis | 38](#)

Learn about open issues in Junos OS Release 22.2R1 for EX Series switches.

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

Forwarding and Sampling

- The fast-lookup-filter with match not supported in FLT Hardware might cause the traffic drop.
[PR1573350](#)

Infrastructure

- On ARM64 platforms such as EX4100, if a live vmcore is attempted to be created, the DUT might get stuck and reboot. [PR1656625](#)

Layer 2 Features

- On QFX5100/EX4600 platforms, if a change related to TPID is made in the device control daemon (dcd), there might be a traffic drop in Packet Forwarding Engine due to failure on L2 learning or interfaces flapping. [PR1477156](#)

Network Management and Monitoring

- A minor memory leak is seen in the event-daemon process when you perform multiple GRES switchovers. [PR1602536](#)

Platform and Infrastructure

- During Routing Engine switchover, interface flap might be seen along with Scheduler slippage.
[PR1541772](#)
- Traffic drops after chassis-control restart when filter is attached and source and destination configurations are enabled. [PR1615548](#)
- Firewall: End-to-End traffic validation fails before applying filter on interface. [PR1634347](#)
- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. [PR1362609](#)
- Runt, fragment and jabber counters are not incrementing on EX4300-MPs. [PR1492605](#)

- Pause frames counters are not getting incremented when pause frames are sent. [PR1580560](#)
- On EX4400 family of devices, sometimes login prompt is not shown after the login session ends. [PR1582754](#)
- EX4400-48MP - VM crash and Virtual Chassis split might be observed with multicast scale scenario. [PR1614145](#)
- When a 100G interface on a QFX5120 is converted to a Virtual Chassis port, the interface stays down as the port is configured as 40G internally. [PR1638156](#)
- With SFP+-10G-CU3M DAC, link remains up on EX4100-48P devices even though the admin is down on peer.

[PR1640799](#)

- Class-of-service buffer-size exact configuration is not supported. The respective configured queue will still use the shared-pool. [PR1644355](#)
- On EX4100 devices, the input pps, bps, and byte counters gets displayed as 0 for some ports while traffic runs without any issues. The interface status gets cleared for 0 to 23 and not cleared for 24 to 47 after interface flaps. [PR1657995](#)
- On EX4100 devices, the Junos telemetry interface FAN and power supply names do not match with CLI. [PR1648739](#)
- Incorrect trap is generated after removal of fan0 in FPC4. [PR1652388](#)
- Momentary traffic loss might be observed with type-5 and NSR/NSB configuration on routing-engine switchover on EX4100 with Junos OS release 22.R1. [PR1655052](#)
- Interop for 1G interfaces between EX4100 SKUs and ACX5448/ACX5448-M/D will not work. [PR1657766](#)
- EX4100 MACsec interface statistics of encrypted/decrypted bytes won't be updating properly after a certain value. [PR1658584](#)
- On EX4600 devices, Virtual-chassis remains in the Unstable state for 3 to 7 minutes, causing traffic loss. [PR1661349](#)
- On EX4600 devices, memory leakage occurs in the eventd process with the longevity test of back to back GRES. [1645852](#)
- Some of interfaces gets zero status in the `monitor interface traffic` command, when traffic gets sent across all interfaces and applied speed of 100m on all 1g copper ports. The issue occurs when you clear the status for interfaces. [PR1661617](#)

- On EX2300 and EX3400 devices, high CPU utilization might be observed when more PoE devices (more than 25 PDs) gets connected to the switch. [PR1667564](#)
- When dot1x enabled access port is member of a secondary pvlan (either community or isolated vlan) by configuration then CLI show dot1x interface detail shows the 'Authenticated Vlan' as primary vlan and 'Authenticated secondary vlan' as configured secondary vlan though client/port is not yet authenticated. Data traffic is still blocked until client is authenticated. This is just an incorrect display and discrepancy with other release. There is no known functional impact. [PR1668144](#)
- On EX4300MP devices, the dcpfe process generates core files when you perform NSSU from Junos OS release 21.1R3.12 to Junos OS release 21.3R3.6. [PR1668414](#)
- On EX4100 devices, traffic does not go through on the management port at link speed 10 and 100M in Junos OS release 22.2R1 and Junos OS release 22.2R1-S1. [PR1676433](#)
- If you insert 1G optic on the uplink ports of EX4100-24mp, EX4100-48p, EX4100-48t, EX4100-24p, and EX4100-24t SKUs, then the activity LED gets lit irrespective of the link Present or Up status. [PR1682633](#)
- On EX4100 devices, secondary USB Type C console port in the front panel does not display proper output. [PR1616315](#)
- On EX4100 devices, Layer 2 IGMP, MLD snooping requires Layer 3 irb to be configured. [PR1681478](#)

Virtual Chassis

- On EX4600 and EX4650 Virtual Chassis (VC), when you reboot the primary Routing Engine, the line card might be disconnected from VC. The Packet Forwarding Engine (PFE) planned restart might be seen on the new backup due to which line card leaves and rejoins the VC automatically. [PR1669241](#)

Resolved Issues

IN THIS SECTION

- Infrastructure | [39](#)
- Interfaces and Chassis | [39](#)
- Layer 2 Ethernet Services | [39](#)

- Platform and Infrastructure | [39](#)
- Routing Protocols | [42](#)

Learn about the issues fixed in this release for EX Series.

Infrastructure

- Recovery snapshot might fail if OAM volume is already mounted. [PR1639991](#)

Interfaces and Chassis

- The vrrpd might crash and generate core files after interface state change. [PR1646480](#)

Layer 2 Ethernet Services

- Option 82 might not be attached on DHCP request packets. [PR1625604](#)
- DHCP packets might not be sent to the clients when forward-only is reconfigured under the routing instance. [PR1651768](#)

Platform and Infrastructure

- EX4300-48MP: Virtual Chassis: NSSU aborted with Backup Routing Engine might be in inconsistent state. [PR1665562](#)
- MAC address learning failure and traffic loss might be observed on VXLAN enabled ports with native-VLAN configured. [PR1663172](#)
- SSH traffic might be affected when filter log action is used. [PR1663126](#)
- Port mirroring traffic is not being flooded on the expected interfaces. [PR1654812](#)

- L2PT might not work for AE interfaces in Q-in-Q environment. [PR1653260](#)
- The inner tag (C-tag) value might get modified to zero for egress traffic when the inner tag values are copied to the outer tag (S-tag). [PR1652976](#)
- L2PT configuration on a transit switch in a Q-in-Q environment breaks L2PT. [PR1650416](#)
- Some interfaces might be down after a power outage or power cycle. [PR1580829](#)
- VSTP might not work in Q in Q environment. [PR1622404](#)
- Traffic loss might be seen when the interface fails to verify the parameter "LOCAL-FAULT." [PR1623215](#)
- The ARP resolution might fail on VRRP enabled interface. [PR1630616](#)
- Application of firewall filters might break connectivity towards the hosts on EX4300. [PR1630935](#)
- The Packet Forwarding Engine might get crash when Virtual Chassis member flaps on EX platforms. [PR1634781](#)
- SCB reset with error : zfchip_scan line = 844 name = failed due to PIO errors [PR1648850](#)
- FXPC might crash and generate a core file due to Segmentation fault during VCCP flap. [PR1655530](#)
- On EX4300 platform, high CPU usage is seen with generation of log message /kernel: %KERN-3: i802_3_slow_recv_input:oam/esmc PDU dropped. [PR1661332](#)
- Junos 'et-' interface stuck and remains down between two particular ports. [PR1535078](#)
- Error message error: syntax error: request-package-validate will be seen on device CLI output during non stop Software Upgrade. [PR1596955](#)
- The device will be unavailable while performing FIPS 140-2/FIPS 140-3 level 2 internal test on FreeBSD 12 based Junos platforms. [PR1623128](#)
- DHCPv6 server binding might not happen when LDRA is enabled along with DHCPv6 snooping. [PR1627600](#)
- System time might not be updated after reboot on EX2300 platform. [PR1627673](#)
- The error message BCM_PVLAN_UTILS:ERR:pfe_bcm_pvlan_utils_get_sec_bd(),789: Failed to get Secondary-bd is logged when a dhcp packet is received on private vlan. [PR1630553](#)
- Unicast ARP packets with the first four bytes of its destination MAC matching to system macs of a transit system get trapped by the system. [PR1632643](#)
- There might be traffic loss for 20 seconds on Virtual Chassis with AE link-protection when rebooting backup FPC. [PR1633115](#)

- The VCPs connected with the AOC cable might not come up after upgrading to Junos OS Release 17.3 or later releases. [PR1633998](#)
- The fxpc process might crash when a MAC is aging out. [PR1634433](#)
- The dot1x ports might be stuck in the connecting state after clearing the dot1x sessions. [PR1634820](#)
- EX4300-48MP - LED state stays OFF in the output of show chassis led for 40G port on PIC 2. [PR1635106](#)
- IPv6 route advertisement sent on management interfaces might cause other devices to fail to get the dhcipv6 address. [PR1635867](#)
- There might be IRB traffic drop when mac-persistence-timer expires. [PR1636422](#)
- Configuring L2PT on a transit switch in a Q-in-Q environment breaks L2PT for other S-VLANs. [PR1637249](#)
- MAC address might not be learned on the new interface after MAC move. [PR1637784](#)
- DHCP snooping table might fail on all Junos platforms to populate MAC address after a VLAN change. [PR1637380](#)
- There might be a delay for the interfaces to come up after reboot/transceiver replacement. [PR1638045](#)
- MAC-move might occur when you configure dhcp-security.. [PR1639926](#)
- The error message dot1xd : devrt_rtsock Don't know how to handle message type 2 is logged even if dot1x is not set. [PR1641304](#)
- The VMcore might be observed on EX platforms in a rare scenario. [PR1641988](#)
- Junos OS: RIB and Packet Forwarding Engines might get out of sync due to a memory leak caused by interface flaps or route churn (CVE-2022-22209). [PR1642172](#)
- There might be traffic impact if you enable persistent-learning on an interface. [PR1643258](#)
- Space issues might be seen on EX3400 devices. [PR1643824](#)
- Traffic loop might occur due to STP ports not created in new master Routing Engine after switchover due to reboot of master Routing Engine on EX4300, EX3400, and EX2300 platforms in Virtual Chassis (VC) scenario. [PR1647000](#)
- An incorrect PEM alarm will be raised on EX4400 devices. [PR1658049](#)

Routing Protocols

- Message Initialize libtask-license first! for mcsnoopd is seen after committing configuration.
[PR1636261](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 42](#)

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

Starting in Junos OS release 21.2R1, all Junos OS products which were previously running on FreeBSD 11.x based Junos OS are migrated to FreeBSD 12.x based Junos OS, except EX4400. Starting with Junos OS release 21.3R1, EX4400 platforms are migrated to FreeBSD 12.x based Junos OS.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

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

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

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

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases.

Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 3: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for JRR Series

IN THIS SECTION

- [What's New | 44](#)
- [What's Changed | 44](#)
- [Known Limitations | 44](#)
- [Open Issues | 44](#)
- [Resolved Issues | 44](#)
- [Migration, Upgrade, and Downgrade Instructions | 45](#)

These release notes accompany Junos OS Release 22.2R1 for the JRR Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

There are no new features or enhancements to existing features in Junos OS Release 22.2R1 for JRR Series Route Reflectors.

What's Changed

There are no changes in behavior and syntax in Junos OS Release 22.2R1 for JRR Series Route Reflectors.

Known Limitations

There are no known limitations in hardware and software in Junos OS 22.2R1 for JRR Series Route Reflectors.

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

Open Issues

There are no known issues in hardware and software in Junos OS Release 22.2R1 for JRR Series Route Reflectors.

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

Resolved Issues

There are no resolved issues in Junos OS Release 22.2R1 for JRR Series Route Reflectors.

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 45](#)

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

For information about software installation and upgrade, see the [JRR200 Route Reflector Quick Start](#) and [Installation and Upgrade Guide](#).

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

We have two types of releases, EOL and EEOL:

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

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

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 4: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for Juniper Secure Connect

IN THIS SECTION

- [What's New | 47](#)
- [What's Changed | 47](#)
- [Known Limitations | 47](#)
- [Open Issues | 47](#)
- [Resolved Issues | 47](#)

These release notes accompany Junos OS Release 22.2R1 for Juniper Secure Connect. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

There are no new features or enhancements to existing features in Junos OS Release 22.2R1 for Juniper Secure Connect.

What's Changed

There are no changes in behavior and syntax in Junos OS Release 22.2R1 for Juniper Secure Connect.

Known Limitations

There are no known limitations in hardware and software in Junos OS 22.2R1 for Juniper Secure Connect.

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

Open Issues

There are no known issues in hardware and software in Junos OS Release 22.2R1 for Juniper Secure Connect.

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

Resolved Issues

There are no resolved issues in Junos OS Release 22.2R1 for Juniper Secure Connect.

Junos OS Release Notes for Junos Fusion for Enterprise

IN THIS SECTION

- [What's New | 48](#)
- [What's Changed | 48](#)
- [Known Limitations | 49](#)
- [Open Issues | 49](#)
- [Resolved Issues | 49](#)
- [Migration, Upgrade, and Downgrade Instructions | 50](#)

These release notes accompany Junos OS Release 22.2R1 for the Junos Fusion for enterprise. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

There are no new features or enhancements to existing features in Junos OS Release 22.2R1 for Junos fusion for enterprise.

What's Changed

There are no changes in behavior and syntax in Junos OS Release 22.2R1 for Junos Fusion for enterprise.

Known Limitations

There are no known limitations in hardware and software in Junos OS 22.2R1 for Junos fusion for enterprise.

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

Open Issues

There are no known issues in hardware and software in Junos OS Release 22.2R1 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

IN THIS SECTION

- [Junos Fusion Enterprise | 49](#)

Learn about the issues fixed in this release 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 Enterprise

- The `sdpd` process generates core file with traces `0x0815e029` in `vfp_cascade_port_discovered,0x0817976d` in `csp_sd_device_discovered` on `b54-rodnik2-sys`. [PR1555597](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device | 50](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 52](#)
- [Preparing the Switch for Satellite Device Conversion | 52](#)
- [Converting a Satellite Device to a Standalone Switch | 54](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 54](#)
- [Downgrading Junos OS | 55](#)

This section contains the procedure to upgrade or downgrade Junos OS and satellite software for a Junos fusion for enterprise. Upgrading or downgrading Junos OS and satellite software might take several hours, depending on the size and configuration of the Junos fusion for enterprise topology.

Basic Procedure for Upgrading Junos OS on an Aggregation Device

When upgrading or downgrading Junos OS for an aggregation device, always use the `junos-install` package. Use other packages (such as the `jbundle` package) only when so instructed by a Juniper Networks support representative. For information about the contents of the `junos-install` package and details of the installation process, see the [Installation and Upgrade Guide](#).



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

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

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

before upgrading or downgrading the routing platform. See the [Junos OS Administration Library](#).

To download and install Junos OS:

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



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

Customers in the United States and Canada, use the following commands, where *n* is the spin number.

```
user@host> request system software add validate reboot source/package-name.n.tgz
```

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

```
user@host> request system software add validate reboot source/package-name.n-limited.tgz
```

Replace *source* with one of the following values:

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

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Preparing the Switch for Satellite Device Conversion

There are multiple methods to upgrade or downgrade satellite software in your Junos fusion for enterprise. See [Configuring or Expanding a Junos fusion for enterprise](#).

For satellite device hardware and software requirements, see [Understanding Junos fusion for enterprise Software and Hardware Requirements](#).

Use the following command to install Junos OS on a switch before converting it into a satellite device:

```
user@host> request system software add validate reboot source/package-name
```



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

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

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

1. Log in to the device using the console port.
2. Clear the device:

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



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

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

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

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

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

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

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

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

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, or preconfiguration. See [Configuring or Expanding a Junos fusion for enterprise](#) for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Switch

If you need to convert a satellite device to a standalone device, you must install a new Junos OS software package on the satellite device and remove it from the Junos fusion topology. For more information, see [Converting a Satellite Device to a Standalone Device](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

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

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

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 5: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Downgrading Junos OS

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



NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

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

Junos OS Release Notes for Junos Fusion for Provider Edge

IN THIS SECTION

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

These release notes accompany Junos OS Release 22.2R1 for Junos Fusion for provider edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

There are no changes in behavior and syntax in Junos OS Release 22.2R1 for Junos Fusion for Enterprise.

What's Changed

There are no changes in behavior and syntax in Junos OS Release 22.2R1 for Junos Fusion for provider edge.

Known Limitations

There are no known limitations in hardware and software in Junos OS 22.2R1 for Junos fusion for provider edge.

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

Open Issues

There are no known issues in hardware and software in Junos OS Release 22.2R1 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

There are no resolved issues in Junos OS Release 22.2R1 for Junos Fusion for provider edge.

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 58](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 60](#)
- [Preparing the Switch for Satellite Device Conversion | 61](#)
- [Converting a Satellite Device to a Standalone Device | 62](#)
- [Upgrading an Aggregation Device | 65](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 65](#)
- [Downgrading from Junos OS Release 22.2 | 66](#)

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 22.2R1 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/jinstall64-22.2R1.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot source/jinstall-22.2R1.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/jinstall64-22.2R1.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot source/jinstall-22.2R1.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 22.2R1 jinstall package, you cannot return to the previously installed software by issuing the `request system software rollback` command. Instead, you must issue the `request system software add validate` command and specify the jinstall package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.

3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite device, you first need to upgrade the target satellite device to an interim Junos OS software version that can be converted to satellite software. For satellite device hardware and software requirements, see [Understanding Junos fusion Software and Hardware Requirements](#)



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

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

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

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

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

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

2. Clear the device:

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



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

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

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

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

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

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

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

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

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, and preconfiguration. See [Configuring Junos fusion for provider edge](#) for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Device

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



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

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

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

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

2. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos fusion** from the drop-down list and select the switch platform series and model for your satellite device.
4. Select the Junos OS Release 14.1X53-D30 software image for your platform.
5. Review and accept the End User License Agreement.
6. Download the software to a local host.
7. Copy the software to the routing platform or to your internal software distribution site.
8. Remove the satellite device from the automatic satellite conversion configuration.

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

[edit]

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

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

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

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

9. Commit the configuration.

To commit the configuration to both Routing Engines:

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

Otherwise, commit the configuration to a single Routing Engine:

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

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

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

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

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

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

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

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

11. Wait for the reboot that accompanies the software installation to complete.
12. When you are prompted to log back into your device, uncable the device from the Junos fusion topology. See [Removing a Transceiver from a QFX Series Device](#) or [Remove a Transceiver](#), as needed. Your device has been removed from Junos fusion.



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

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 22.2R1, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

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

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

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases.

Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 6: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Downgrading from Junos OS Release 22.2

To downgrade from Release 22.2 to another supported release, follow the procedure for upgrading, but replace the 22.2 `jinstall` package with one that corresponds to the appropriate release.



NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for MX Series

IN THIS SECTION

- [What's New | 67](#)

- [What's Changed | 93](#)
- [Known Limitations | 100](#)
- [Open Issues | 102](#)
- [Resolved Issues | 109](#)
- [Migration, Upgrade, and Downgrade Instructions | 125](#)

These release notes accompany Junos OS Release 22.2R1 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 22.2R1-S2 | 67](#)
- [What's New in 22.2R1 | 79](#)

Learn about new features introduced in this release for the MX Series routers.

What's New in 22.2R1-S2

IN THIS SECTION

- [Hardware | 68](#)
- [Software Installation and Upgrade | 78](#)
- [User Authentication | 78](#)

Learn about new features or enhancements to existing features in Junos OS Release 22.2R1-S2 for the MX Series routers.

Hardware

- **New Routing Engine RE-S-X6-128G-K with TPM 2.0 (MX240, MX480, and MX960)**—In Junos OS Release 22.2R1S2, we introduce the RE-S-X6-128G-K, a new Routing Engine integrated with Trusted Platform Module 2.0 (TPM 2.0). This new Routing Engine is an upgrade to the existing Routing Engine RE-S-X6-128G-S.



NOTE: The RE-S-X6-128G-K Routing Engine must be used with either SCBE2-MX or SCBE3-MX.

The key features of the RE-S-X6-128G-K include:

- Digital cryptographic identity (also called device ID or DevID) embedded in TP M2.0
- RFC 8572-based secure zero-touch provisioning (secure ZTP)

[See [RE-S-X6-128G-K Routing Engine Description](#).]

- **New MX304 Universal Routing Platform**—Starting in Junos OS Release 22.2R1-S2, we introduce the MX304 router—a 2-U, compact modular system that can scale up to 4.8-Tbps capacity. This bandwidth gives hyperscalers, cloud providers, and service providers the performance and scalability needed as networks grow. The router supports 400GbE, 100GbE, 50GbE, 40GbE, 25GbE, and 10GbE interfaces. It has pluggable Routing Engines (it supports one or two Routing Engines), redundant power, and cooling capability. It accepts up to three line-card MICs (LMICs). Each LMIC has 1 YT chip and 1.6 Tbps of forwarding capacity. It supports 4x400-Gbps ports, 16x100-Gbps ports, or a combination.

Table 7: Features Supported on MX304

Feature	Description
Chassis	<ul style="list-style-type: none"> Fabric management support includes fabric hardening, fabric board control, and fault handling. Fabric management includes support for built-in SFB and line-card MIC (LMIC model number JNP304-LMIC16-BASE). MX304 routers support three LMICs (additional LMIC model number MX304-LMIC16-BASE). <p>The SFB provides 18 fabric links to each PFE. There is no SFB fabric redundancy support.</p> <p>[See Fabric Plane Management.]</p> <ul style="list-style-type: none"> Limited-encryption Junos OS image and boot restriction <p>[See Junos OS Editions.]</p> <ul style="list-style-type: none"> Support for platform resiliency <p>[See show system errors active.]</p>
Class of service (CoS)	<ul style="list-style-type: none"> Forwarding CoS and hierarchical CoS (HCoS) support. <p>[See Understanding Class of Service and Hierarchical Class of Service for Subscriber Management Overview.]</p>
Distributed denial-of service (DDoS)	<ul style="list-style-type: none"> DDoS protection is enabled by default. <p>[See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]</p>

Table 7: Features Supported on MX304 (*Continued*)

Feature	Description
Flow monitoring	<ul style="list-style-type: none"> Support for Inline services—We support the following Inline services: <ul style="list-style-type: none"> Inline active flow monitoring Inline monitoring Video monitoring FlowTapLite <p>[See Monitoring, Sampling, and Collection Services Interfaces User Guide.]</p> <ul style="list-style-type: none"> Support for Routing-Engine-based traffic sampling You can configure Routing-Engine-based traffic sampling. Traffic sampling enables you to copy traffic to a line card that performs flow accounting while the router forwards the packet to its original destination. You configure either an input or an output firewall filter with a matching term that contains the <code>then sample</code> statement. Routing-Engine-based traffic sampling supports only the version 5 and version 8 formats for exporting flow records. <p>[See Configuring Traffic Sampling on MX, M and T Series Routers.]</p>
Hardware	<ul style="list-style-type: none"> The MX304 router contains pluggable Routing Engines and supports up to three LMICs. Each LMIC supports 4x400-Gbps ports, 16x100-Gbps ports, or a combination. The MX304 router has two dedicated AC, DC, or HVAC/HVDC power supply modules and front-to-back cooling. <p>MX304 Universal Routing Platform Hardware Guide</p> <ul style="list-style-type: none"> Supported transceivers, optical interfaces, and DAC cables—Select your product in the Hardware Compatibility Tool to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

Table 7: Features Supported on MX304 (*Continued*)

Feature	Description
High availability (HA) and resiliency	<ul style="list-style-type: none"> Support for BFD: <ul style="list-style-type: none"> Centralized, distributed, inline, single-hop, multihop, and micro-BFD. BFD over integrated routing and bridging (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). <p>[See Understanding BFD for Static Routes for Faster Network Failure Detection, and Bidirectional Forwarding Detection (BFD).]</p> <ul style="list-style-type: none"> Resiliency support for Packet Forwarding Engine and the built-in Switch Fabric Board (SFB). <p>[See show system errors active.]</p>

Table 7: Features Supported on MX304 (*Continued*)

Feature	Description
Interfaces	<ul style="list-style-type: none"> MX304 introduces a pluggable 4x400GbE and 16x100GbE Combo LMIC. MX304 can deliver a bandwidth of up to 4.8Tbps. Each MX304 LMIC hosts two Packet Forwarding Engines with overall bandwidth of 1.6 Tbps. Each PFE is capable of 800G and overall it becomes 1.6 Tbps. <p>Each port supports 10-Gbps, 25-Gbps, 40-Gbps, 50-Gbps, 100-Gbps, 200-Gbps, and 400-Gbps interface speeds using different optics.</p> <p>You can channelize the interfaces as follows:</p> <ul style="list-style-type: none"> Four 10 GbE interfaces Four 25 GbE interfaces One 100 GbE interfaces Two 100 GbE interfaces Four 100 GbE interfaces <p>Note that we support 40G channelization on all odd ports, but alternate ports should be empty.</p> <p>You can configure the port speed at the [edit chassis] hierarchy level.</p> <p>[See Port Speed.]</p> <ul style="list-style-type: none"> Supports transceivers, optical interfaces, and direct attach copper (DAC) cables on MX304. <p>[See Hardware Compatibility Tool , and optics-options.]</p> <ul style="list-style-type: none"> Support for flexible tunnel interfaces <p>[See Flexible Tunnel Interfaces Overview.]</p>

Table 7: Features Supported on MX304 (Continued)

Feature	Description
Juniper telemetryinterface (JTI)	<ul style="list-style-type: none"> • NPU and CPU memory utilization telemetry sensor support in JTI—You can use JTI to stream network processing unit (NPU) and CPU statistics to an outside collector from an MX304 router. Include the following sensors in a remote procedure calls (gRPC) or gRPC network management interface (gNMI) subscription: <ul style="list-style-type: none"> • <code>/junos/system/linecard/cpu/memory/</code> • <code>/junos/system/linecard/npu/memory/</code> • <code>/junos/system/linecard/npu/utilization/</code> [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).] • Logical interface statistics for IPv4 and IPv6 family counters—You can stream per-family logical interface input and output counters for IPv4 and IPv6 traffic using JTI and gRPC to an outside collector. <ul style="list-style-type: none"> Include the resource paths <code>/junos/system/linecard/interface/</code> <code>logical/family/ipv4/usage/</code> and <code>/junos/system/linecard/</code> <code>interface/logical/family/ipv6/usage/</code> in a gRPC subscription. [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).] • Transceiver diagnostics sensor support in JTI—JTI supports the OpenConfig transceiver model <code>openconfig-platform-transceiver.yang</code> 0.5.0. You can deliver ON_CHANGE transceiver statistics to an outside collector using remote procedure calls (gRPC) or gRPC network management interface (gNMI) services. <ul style="list-style-type: none"> [See Telemetry Sensor Explorer.]

Table 7: Features Supported on MX304 (Continued)

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> Support for Layer 2 features <p>[See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation, Understanding Layer 2 Bridge Domains, Understanding Layer 2 Learning and Forwarding, and Introduction to OAM Connectivity Fault Management (CFM).]</p> <ul style="list-style-type: none"> Support for Layer2 Ethernet services over GRE tunnel interfaces <p>[See Configuring Layer 2 Ethernet Services over GRE Tunnel Interfaces.]</p>
Layer 3 features	<ul style="list-style-type: none"> Support for Layer 3 features <p>[See MPLS Overview, Multicast Overview, Tunnel Services Overview, and Understanding Next-Generation MVPN Control Plane.]</p> <ul style="list-style-type: none"> Load balancing support: <ul style="list-style-type: none"> Enhanced hash key options. Consistent flow hashing, source IP-only hashing, and destination IP-only hashing. Symmetrical load balancing over 802.3 and LAGs. <p>[See Understanding Per-Packet Load Balancing.]</p>

Table 7: Features Supported on MX304 (*Continued*)

Feature	Description
Layer 3 VPN	<ul style="list-style-type: none"> Anti-spoofing protection for next-hop-based dynamic tunnelsWe've added antispoofing capabilities to IPv4 tunnels and IPv4 data traffic. Antispoofing for next-hop-based dynamic tunnels can detect and prevent a compromised virtual machine (inner source reverse path forwarding check) but does not apply to a compromised server that is label-spoofing. The antispoofing protection is effective when the VRF routing instance has label-switched interfaces (LSIs) using vrf-table-label or virtual tunnel (VT) interfaces. We do not support antispoofing protection for per-next-hop labels on VRF routing instances. <p>[See Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels Overview and Example: Configuring Anti-Spoofing Protection for Next-Hop-Based Dynamic Tunnels.]</p>
MACsec	<ul style="list-style-type: none"> Support for Media Access Control Security (MACsec), including AES-256 encryption, extended packet numbering, and fail-open mode <p>[See Configuring Media Access Control Security (MACsec) on Routers.]</p> <ul style="list-style-type: none"> MACsec bounded delay protection <p>[See bounded-delay.]</p>

Table 7: Features Supported on MX304 (Continued)

Feature	Description
Multicast	<ul style="list-style-type: none"> Auto LSP Policer support: <ul style="list-style-type: none"> Multicast load balancing of point-to-multipoint (P2MP) label-switched-paths (LSPs) over aggregated Ethernet child links. Automatic policers for MPLS P2MP LSPs. Display of packet and byte statistics for sub-LSPs of a P2MP LSP. GRES and graceful restart for MPLS P2MP LSPs. Multicast virtual private network (MVPN) extranet or overlapping functionality. <p>[See Example: Configuring Multicast Load Balancing over Aggregated Ethernet Links, and Point-to-Multipoint LSP Configuration]</p>
Network management and monitoring	<ul style="list-style-type: none"> Support for port mirroring <p>[See Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers.]</p> <ul style="list-style-type: none"> Support for configuring ITU-T Y.1731 standard-compliant Ethernet synthetic loss measurement (ETH-SLM) and Ethernet delay measurement (ETH- DM) capabilities <p>[See ITU-T Y.1731 Ethernet Service OAM Overview.]</p>
Routing policy and firewall filters	<ul style="list-style-type: none"> Support for forwarding firewalls <p>[See Understanding Firewall Filter Match Conditions, Overview of Policers, Fast Update Filters Overview, Service Filter Overview, and Understanding Firewall Filter Fast Lookup Filter.]</p>

Table 7: Features Supported on MX304 (*Continued*)

Feature	Description
Services applications	<ul style="list-style-type: none"> • Inline Services support: <ul style="list-style-type: none"> • Inline NAT—NAT44 and NPTv6 • Inline softwires—Mapping of Address and Port with Encapsulation (MAP-E) and IPv6 rapid deployment (6rd) • Inline J-Flow • Inline monitoring • Video monitoring • FlowTapLite <p>[See Inline NAT, Configuring Mapping of Address and Port with Encapsulation (MAP-E), Configuring Inline 6rd, and Monitoring, Sampling, and Collection Services Interfaces User Guide.]</p> • Support for RFC 2544-based benchmarking tests <p>[See Understanding RFC2544-Based Benchmarking Tests on MX Series Routers.]</p> • Support for Two-Way Active Measurement Protocol (TWAMP) and Real-Time Performance Monitoring (RPM) <p>[See Understand Two-Way Active Measurement Protocol, and Real-Time Performance Monitoring.]</p> • DHCP security—The MX304 router supports the following DHCP security features: <ul style="list-style-type: none"> • DHCP snooping with Option 82. • DHCPv6 snooping with Option 16, Option 18, Option 37, and Option 79. • Lightweight DHCPv6 relay agent. <p>[See DHCP Snooping.]</p>

Table 7: Features Supported on MX304 (*Continued*)

Feature	Description
Software installation and upgrade	<ul style="list-style-type: none"> Support for secure boot [See Secure Boot.] Support for zero-touch provisioning (ZTP) on the management interface. ZTP automates the provisioning of the device configuration and software upgrade over the management interface of the Routing Engine. [See Zero Touch Provisioning Overview.]

Software Installation and Upgrade

- Secure Zero Touch Provisioning (SZTP) (MX240, MX480, and MX960)**—Starting in Junos OS Release 22.2R1S2, you can use RFC-8572-based SZTP to bootstrap your remotely located network devices that are in a factory-default state. SZTP enables mutual authentication between the bootstrap server and the network device before the remote network device is accessed for initiating zero touch provisioning.

To enable mutual authentication, you need a unique digital voucher, which is generated based on the DevID (Digital Device ID or Cryptographic Digital Identity) of the network device. The DevID is embedded inside the Trusted Platform Module (TPM) 2.0 chip on the network device. Juniper Networks issues a digital voucher to customers for each eligible network device.

[See [Secure Zero Touch Provisioning](#) and [Generate Voucher Certificate](#).]

User Authentication

- Support for File-system Encryption with Trusted Platform Module (TPM 2.0) (MX240, MX480, MX960)**—Starting in Junos OS Release 22.2R1S2, you can encrypt file-system data residing on your MX Series device hard disk drives with Trusted Platform Module 2.0 (TPM 2.0). TPM is a chip used for the identification and authentication of a device on the network and to ensure the software loaded on the system is in the correct state when it started up.

Limitations

- Installing the old image previous to 22.2R1 and performing an ISSU when encryption is enabled on the disk does not issue a warning or an error.

- Users with root access permission can only perform file-system encryption on the device.
- Hard disk encryption is not automatically applicable on the newly inserted drive.
- File-system encryption is applicable only on these MX Series Routing Engines: RE-S-X6-128G-K-BB, RE-S-X6-128G-K-R, and RE-S-X6-128G-K-S.
- The `request system filesystem encryption keys delete` command deletes disk keys on lock and unlock drives. Locked drives are due to PCR mismatch or disk inserted from another system.
- We do not support enabling encryption with GRES. You must enable encryption on each Routing Engine.
- Forced clearing of TPM keys results in an unexpected behavior.
- You cannot use automatic recovery feature after deleting keyslots.
- We do not support VM host snapshot recovery.
- The `show system filesystem encryption status` command display information about the specific Routing Engine only.

[See [Encryption with TPM](#).]

What's New in 22.2R1

IN THIS SECTION

- [EVPN | 80](#)
- [High Availability | 81](#)
- [Interfaces | 81](#)
- [IP Tunneling | 82](#)
- [Junos Telemetry Interface | 82](#)
- [Licensing | 83](#)
- [MACsec | 84](#)
- [MPLS | 84](#)
- [Platform and Infrastructure | 85](#)
- [Precision Time Protocol \(PTP\) | 85](#)
- [Routing Policy and Firewall Filters | 86](#)
- [Routing Protocols | 86](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 89](#)

- Software Installation and Upgrade | [89](#)
- Subscriber Management and Services | [90](#)
- VPNs | [91](#)
- Additional Features | [91](#)

Learn about new features or enhancements to existing features in Junos OS Release 22.2R1 for the MX Series routers.

EVPN

- **Support for blocking asymmetric EVPN Type 5 routes (MX960, QFX5110, and QFX10002)**—Starting in Junos OS Release 22.2R1, you can configure the local node to reject asymmetric EVPN Type 5 routes on EVPN-VXLAN networks. The local node examines the incoming EVPN Type 5 route packets and rejects the route when the virtual network identifier (VNI) in the ingress route differs from the locally configured VNI.

To block asymmetric EVPN Type 5 routes, include the `reject-asymmetric-vni` statement at the `[edit routing-instance routing-instance-name protocols evpn ip-prefix-routes]` hierarchy level.

[See [EVPN Type 5 Route with VXLAN encapsulation for EVPN-VXLAN](#) and [ip-prefix-routes](#).]

- **Automatically derived ESI configuration (MX Series, QFX5100, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX10002, QFX10002-60C, QFX10008, and QFX10016)**—In the current implementation, Junos OS derives the Ethernet segment identifier (ESI) from the system ID and the administrative key on the local multihomed provider edge (PE) device that is a part of the LACP link (actor). Starting in Junos OS Release 22.2R1, you can also configure the multihomed devices on an EVPN-VXLAN network to automatically generate the ESI from:

- The system ID and administrative key on the remote customer edge (CE) device (partner).
- The locally configured `mac` and local discriminator values.

To automatically derive the ESI using the system ID and administrative key on the remote CE device, include `type-1-lacp` at the `[edit interfaces aeX aggregated-ether-options lacp auto-derive]` hierarchy level.

To automatically derive the ESI using locally configured values, configure `mac` and `local-discriminator` at the `[edit interfaces aeX aggregated-ether-options lacp auto-derive type-3-system-mac]` hierarchy level.

[See [Understanding Automatically Generated ESIs in EVPN Networks](#).]

- **EVPN active/active redundancy, aliasing, and mass MAC withdrawal (MX Series and vMX)**—Starting in Junos OS Release 22.2R1, the listed devices support EVPN active/active redundancy, aliasing, and

mass MAC withdrawal, integrated with VXLAN in the data plane. These features provide resilient inter-data center connectivity to the established Data Center Interconnect (DCI) technologies. This new support builds an end-to-end DCI solution by integrating EVPN active/active multicast with DP VXLAN.

Use existing configuration statements to configure active/active redundancy at the ESI level on the loopback (lo0) interface. Include lo0 as the virtual tunnel endpoint (VTEP) interface in the routing instance.

[See [EVPN-over-VXLAN Supported Functionality](#).]

- **Support for BGP domain path attribute in EVPN Type 5 and IPVPN routes on gateway provider edge (PE) devices (MX480, MX960, and vMX)**—Starting in Junos OS Release 22.2R1, you can configure the BGP D-PATH attribute on your gateway PE device to add a domain ID to BGP routes. The BGP D-PATH attribute enables gateway PE devices to identify the domains through which EVPN IP prefix routes and IP-VPN routes have traversed. Additionally, the BGP D-PATH attribute uses its path selection algorithm to install the best routes in gateway PE device IP virtual routing and forwarding (VRF) tables, which prevents routes from looping.

To configure the BGP D-PATH attribute on all of your configured virtual routing and forwarding instances on the gateway PE device, enable the `uniform-propagation-mode` statement with the `domain-id` option in the `[edit routing-instances]` hierarchy. When you configure the statement, it is also enabled at the global `[edit routing-options uniform-propagation-mode domain-id <type>]` hierarchy level. Use the `<type>` variable to specify the type of Inter-Subnet Forwarding (ISF) and Subsequent Address Family Identifiers (SAFIs) you use to advertise IP prefix routes.

[See [uniform-propagation-mode](#).]

High Availability

- **Support for OSPFv3 and IS-IS BFD sessions using IPv6 link-local addresses (MX240, MX480, MX960, MX2010, MX2020, and MX10008 devices with MPC10 and MPC11 line cards)**—Starting in Junos OS 22.2R1, MPC10 and MPC11 line cards on your network devices support OSPFv3 and IS-IS BFD sessions that use IPv6 link-local addresses.

[See [No Link Title](#).]

Interfaces

- **Support for 400G ZR DWDM optics (MX2010, MX2020 and MX10K-LC9600)**—Starting in Junos OS Release 22.2R1, we support 400G ZR DWDM optics. We support the following 400G ZR features:
 - 75-GHz grid spacing.
 - Optical loopbacks. To enable an optical loopback, use the existing `loopback` statement at the `[edit interfaces optics-options]` hierarchy level.

- Performance monitoring and Threshold-crossing alert (TCA). You can view the current and historical performance monitoring metrics, which are accumulated into 15-minute and 1-day interval bins, by using the `show interfaces transport pm` command. You can thus manage optical transport link efficiently. TCAs provide the management system an early indication of the deteriorating health of an optical network connection when the performance parameter that you monitor crosses a certain threshold.

[See [optics-options](#), [show interfaces](#), [show interfaces diagnostics optics](#), [show chassis pic](#), [show chassis hardware](#), and [show interfaces transport pm](#).]

- **MPC7 and MPC5 support with SFB3 (MX2010 and MX2020)**—Starting in Junos OS Release 22.2R1, we support MPC7 and MPC5 with the SFB3 Switch Fabric Board. The SFB3 interoperates with MPC11.

[See [MX2000 Switch Fabric Board \(SFB\) Overview](#).]

- **MPC2E-NG and MPC3E-NG support with SFB3 (MX2010 and MX2020)**—Starting in Junos OS Release 22.2R1, we support MPC2E-NG (with ethernet MICs only) and MPC3E-NG (with ethernet MICs only) with the SFB3 Switch Fabric Board. The SFB3 interoperates with MPC11. This feature is supported only on Ethernet MICs.

[See [MX2000 Switch Fabric Board \(SFB\) Overview](#) .]

IP Tunneling

- **Sharding support for dynamic IP-over-IP tunneling (MX240, MX480, MX960, PTX1000, PTX10001, and QFX10002)**—Starting in Junos OS Release 22.2R1, we are supporting sharding for dynamic tunnels that are created as a result of BGP route resolution over a tunnel route. BGP uses this tunnel route as a helper route for route resolution.

Junos Telemetry Interface

- **Support for breakout port state sensor (MX240, MX480, MX960, MX2010, MX2020, MX10008, MX10016, and PTX5000)**—Junos OS Release 22.2R1 introduces the breakout port state sensor / `components/component/port/breakout-mode/groups/group/state` for Junos telemetry interface (JTI) based on the OpenConfig data model `openconfig-platform-port.yang` version 0.4.0. This sensor provides the operational state data for the breakout group identified by the index on platforms running on JUNOS OS.

[See [Telemetry Sensor Explorer](#) and [sensor \(Junos Telemetry Interface\)](#).]

- **Support for CPU state sensor (ACX710, ACX5448, MX204, MX240, MX150, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, and PTX10002)**—Starting in Junos OS Release 22.2R1, use the resource path `/system/cpus/cpu/state/` to export CPU

parameters and including CPU usage per process and CPU usage per Routing Engine core information from a device to a collector.

[See [Telemetry Sensor Explorer](#).]

- **Support for forwarding table sensor (MX2020 and PTX5000)**—Junos OS Release 22.2R1 extends support for forwarding information base (FIB) streaming on JTI to include non-default virtual routing and forwarding (VRF) instances.

[See [Telemetry Sensor Explorer](#).]

- **Support for Ethernet interface sensors (MX240, MX480, MX960, MX2010, MX2020, MX10003, MX10008, MX10016, and PTX5000)**—Starting in Junos OS Release 22.2R1, use the subscription path `/interfaces/interface/` or `/interfaces/interface/ethernet/state` to stream Ethernet packet statistics from a device to a collector.

[See [Telemetry Sensor Explorer](#).]

- **Network instance support enhancements (ACX710, ACX5448, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, and PTX10002)**—Starting in Junos OS Release 22.2R1, JTI supports new sensors for network instance statistics for the OpenConfig modules `openconfig-network-instance.yang` and `openconfig-routing-policy.yang`. The support includes OpenConfig configuration and streaming of state data.

[See [Telemetry Sensor Explorer](#) for telemetry support and [OpenConfig User Guide](#) for configuration.]

- **Support for QoS sensor (MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, and vMX)**—Junos OS Release 22.2R1 introduces QoS sensors for JTI based on the OpenConfig data model `open-config-qos` version 0.3.0.

[See [Telemetry Sensor Explorer](#).]

- **Junos node slicing support for platform sensors (MX480, MX960, MX2010, MX2020)**—Junos OS Release 22.2R1 extends support for platform streaming on JTI to include Junos node slicing environments.

Licensing

- **Juniper Agile Licensing (MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10008, and MX10016)**—Starting in Junos OS Release 22.2R1, the listed MX Series devices support Juniper Agile Licensing.

Juniper Agile Licensing provides simplified and centralized license administration and deployment. You can use Juniper Agile Licensing to install and manage licenses for hardware and software features.

[See [Flex Software License for MX](#) and [Managing Licenses](#).]

MACsec

- **Certificate-based authentication and encryption for MACsec (MX Series)**—Starting in Junos OS Release 22.2R1, you can enable MACsec on links connecting switches or routers using certificate-based authentication and encryption. Connected devices can mutually authenticate using 802.1X over Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) and dynamically derive the connectivity association key (CAK) for encryption.

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

MPLS

- **Support for RSVP delay constraint (cRPD, MX960, and PTX10008)**—Starting in Junos OS Release 22.2R1, you can configure RSVP label-switched paths (LSPs) to use a delay metric for computing the path. To configure, use the new CLI options that we've introduced under the [edit protocols mpls label-switched-path *name*] hierarchy. We've also updated the outputs of the following show commands:
 - show ted link detail
 - show ted database extensive
 - show route protocol bgp table lsdist.0 extensive
 - show spring-traffic-engineering lsp detail
 - show express-segments name *name* detail
 - show mpls lsp detail
- **Support for ingress and transit chained CNHs for BGP Labeled Unicast (BGP-LU) IPv4 (MX204, MX480, MX960, MX10003, and vMX)**—Starting in Junos OS Release 22.2R1, you can configure the chained composite next hops (CNHs) for devices handling ingress or transit traffic in the network. We've added support only for the following options on the listed MX Series devices:
 - BGP-LU for IPv4 on the ingress router—set routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp inet
 - BGP-LU on the transit router—set routing-options forwarding-table chained-composite-next-hop transit labeled-bgp

You can also configure class of service (CoS) and define rewrite rules for ingress and transit chained CNHs for BGP-LU.



NOTE: This feature is supported only on MPC9 and the previous models of line cards.

- **CBF support for colored SR-TE (MX Series)**— Starting in Junos OS 22.2R1 Release, you can create a multipath route between different transport tunnels in any inet6.3 RIBs to resolve service routes. We also support storage of transport-class in nexthop gateways, support preserve nh hierarchy for policy multipath, and extend cos policy to support transport-class color so that selective traffic can be steered over one of SR-TE or IGP flex-algo transport tunnels.

To match transport-class key in FRR indirect keys structure, use transport-class color <color> at the [edit class-of-service forwarding-policy next-hop-map *next-hop-map-name* forwarding-class *forwarding-class-name*] hierarchy level.

To store a color value as FRR indirect keys structure for the leaked inet6.3 routes, use set best-effort color *color* at the [edit routing-options transport-class] hierarchy level.

To enable expanded nh hierarchy support for policy-multipath routes for RIBs, use set preserve-nexthop-hierarchy at the [edit routing-options rib *rib-name* policy-multipath] hierarchy level. For any inet[6].3 RIBs, the preserve-nexthop-hierarchy is enabled by default.

[See [forwarding-class \(Forwarding Policy\)](#), [policy-multipath](#), and [transport-class](#).]

Platform and Infrastructure

- **Support for monitoring link degradation (MX10008 and MX10016 with MX10K-LC480)**—Starting in Junos OS Release 22.2R1, you can monitor link degradation on the MX10K-LC480 line card.

See [[Link Degrade Monitoring Overview](#).]

- **Reset the Packet Forwarding Engine (MX10008 and MX10016 with MX10K-LC480)**—Starting in Junos OS Release 22.2R1, you can reset the Packet Forwarding Engine on the MX10K-LC480 line card.

See [[show chassis fpc errors](#).]

Precision Time Protocol (PTP)

- **G.8275.1 profile with BITS as a frequency source in hybrid mode (MX10008 with JNP10008 SFB and MX10K-LC2101 line cards)**—Starting in Junos OS Release 22.2R1, you can configure Building Integrated Timing Supply (BITS) as a frequency source with the G.8275.1 profile in PTP hybrid mode. G.8275.1 also supports PTPoE over LAG with BITS as a frequency source.

If you configure both Synchronous Ethernet and BITS as the frequency source, then based on the clock selection, the device chooses either Synchronous Ethernet or BITS as the frequency source in the hybrid mode.

[See [show ptp hybrid](#) and [show chassis synchronization \(MX Series Router\)](#).]

Routing Policy and Firewall Filters

- **Network slicing (MX480, MX960, MX2020, and MX10003)**— Starting in Junos OS Release 22.2R1, you can provision network slices using a combination of CoS and firewall filter configuration. You can use new CoS configuration statements to create slice-based hierarchical queues under any physical interface. We've introduced a new firewall filter matching condition and action named `slice` for family `inet`, `inet6`, and any filters to capture and mark matched packets from and to slices. `slice` is also a new routing policy action to mark packets that match routes with the slice identifier. You can use a new routing policy action named `filter` to bind a named family any filter to the next hop of the routes that match the route policy. To view slice statistics, use the following CLI command:

- `show route extensive expanded-nh` to view the slice and filter information bound to the next hop.
- **Filter based on 6-tuple lookup in inner GTP encapsulated packet (MX240, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10008, and MX10016)**—

Starting in Junos OS Release 22.2R1, Junos OS on the listed MX Series devices supports filter match on the GPRS header (TEID, Version) and inner IP header (5 tuples: Source IP, Destination IP, Source Port, Destination Port, Protocol) in the GTP-C packet.

- **Support for multiple named validation databases from multiple sources (MX204 and PTX10016)**— Starting in Junos OS Release 22.2R1, we support multiple named validation databases from multiple sources. You can also consult validation databases across instances and track RIBs that consult the various databases to enable notification when entries are modified.

To Specify a named route-validation database, use `validation-state (invalid | valid)` option at the [`edit routing-options validation database <database-name> static record <destination> maximum-length <prefix-length> origin-autonomous-system <as-number>`] hierarchy level.

To Specify target route-validation database for a validation session, use `database <database-name>` option at the [`edit routing-options validation group <group-name> session`] hierarchy level.

To specify validation database, use `validation-database-instance` option at the [`edit policy-statement <policy-name> term <term-name> from`] hierarchy level.

[See [policy-statement](#), [session \(Origin Validation for BGP\)](#), and [validation \(Origin Validation for BGP\)](#).]

Routing Protocols

- **DCSPF support for SR-TE with Flex Algo (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, vMX, PTX1000, PTX3000, and PTX5000)**—Starting in Junos OS Release 22.2R1, we support the flexible algorithm (Flex Algo) as a constraint in the compute profile of a segment routing-traffic engineering (SR-TE) LSP. The computation combines any constraints in the compute profile with the ones in the Flex Algo definition to find the resultant path. It uses the Flex Algo segment identifiers (SIDs) in the configuration to compress the resultant path.

We support the feature only for IPv4 SR-MPLS SIDs. You can use SR-TE policy constraints to further fine-tune Flex Algo constraints.

[See [Enabling Distributed CSPF for Segment Routing LSPs.](#)]

- **TCP-AO for RPKI validation sessions (MX204, MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2008, MX2010, MX2020, PTX1000, PTX10002, PTX10008, PTX10016, and vRR) –** Starting in Junos OS Release 22.2R1, you can use TCP Authentication Option (TCP-AO) to authenticate resource public key infrastructure (RPKI) validation sessions for securing the Internet's routing infrastructure, such as BGP. Using RPKI, legitimate holders of Internet number resources can control the operation of Internet routing protocols to prevent route hijacking and other attacks.

To enable a TCP-AO chain to authenticate an RPKI validation session, use `authentication-algorithm ao` and the configured authentication-key-chain *keychain* at the [edit routing-options validation group *group_name* session *address* and [edit routing-options validation group *group_name* hierarchy levels.

See [\[TCP Authentication Option \(TCP-AO\).\]](#)

- **Nonstop active routing (NSR) support with BGP RIP sharding and BGP UpdateIO features (ACX5048, ACX5096, ACX5448, MX240, MX960, MX2008, MX10016, and PTX5000) –** Starting in Junos OS Release 22.2R1, we've enabled nonstop routing (NSR) for BGP RIP sharding and BGP UpdateIO features. With NSR enabled, the backup Routing Engine and backup routing protocol process (rpd) become the primary Routing Engine without negatively affecting the BGP peering sessions with the neighbors if the primary Routing Engine fails. The backup rpd processes the replicated BGP control-plane information and populates the route state in the same multithreaded manner as in the primary rpd.

After you configure NSR, the `show bgp neighbor` and `show bgp summary` commands display the information about the specific shards in the backup Routing Engine. To display the replicated information for a specific shard in the `show bgp replication` command, use the `rib-sharding shard-name` option.

See [\[show bgp neighbor, show bgp summary, show bgp replication, and BGP Overview.\]](#)

- **Layer 3 VPN service interworking between SRv6 and MPLS (MX10008) –** Junos OS Release 22.2R1 supports the Layer 3 VPN service between segment routing over IPv6 (SRv6) and MPLS. The incremental deployment of SRv6 into existing networks requires SRv6 to interwork and coexist with MPLS (RSVP).

See [\[How to Enable SRv6 Network Programming and Layer 3 VPN Services over SRv6 in BGP Networks.\]](#)

- **SRv6 support for multi-instance IS-IS (MX10008) –** Starting in Junos OS Release 22.2R1, we support SRv6 across multi-instance IS-IS. Using this feature, you can run SRv6 on multiple independent IS-IS instances simultaneously in the same router. These multiple instances will act same as the default standard IS-IS instance. The following SID functionality is supported in default or standard IS-IS instance. The performance of enabling SRv6 on multi-instance IS-IS is similar to the one with the SRv6-enabled standard IS-IS IGP instance.

See [[How to Configure Multiple Independent IGP Instances of IS-IS,isis-instance](#).]

- **Support for link delay measurement using TWAMP light and advertising in ISIS (MX240, MX480, MX960, MX2010, MX2020, and MX10008)**—Starting in Junos OS Release 22.2R1, you can measure link delay using TWAMP light and advertise various performance metrics in IP networks using IS-IS. You can use the IS-IS metrics to make path-selection decisions based on network performance.

[See [How to Enable Link Delay Measurement and Advertising in IS-IS](#).]

- **BGP extended route retention (MX960, PTX1000, and QFX10002)**—In Junos OS Release 22.2R1, we've enhanced the long-lived graceful restart (LLGR) capabilities for a BGP helper device. With this feature enabled, Junos OS supports LLGR helper mode regardless of the BGP peer LLGR capabilities. We've introduced a new configuration statement `extended-route-retention` at the [edit protocols bgp group neighbor graceful-restart long-lived] hierarchy level. We've also updated the outputs of the following operational commands:

- `show bgp neighbor`
- `show route extensive`

[See [graceful-restart-long-lived-edit-protocols-bgp](#).]

- **Anomaly checker for rpd object reference count (MX Series, PTX Series, and QFX Series)**—In Junos OS Release 22.2R1, we introduce a generic reference count infrastructure that all the modules in rpd can use. The module maintains lock and unlock statistics corresponding to each object type in use. Any application can call the `refcount increment` or `decrement` API when an object is referred. The module also provides a mechanism to detect anomalies such as a leak or overflow in an object's `refcount`.
- **Origin validation communities conversion to keywords (MX10008 and PTX10016)**—Starting in Junos OS Release 22.2R1, you can choose to accept or reject the origin validation extended communities received from an eBGP peer. The default behavior of Origin Validation State Extended Community (OVS EC) changes to `rejected` if the extended community is received from an eBGP peer. You can configure your device to accept the community when needed. We also support the configuration of distinguished communities with keywords (valid, invalid, and unknown) at all the three layers of the BGP configuration hierarchy—global, group, and per-neighbor. If you enable the OVS EC at a hierarchy level, it's enabled for the lower levels as well. However, you can choose to disable it explicitly at a lower layer if required at any instance.
- **BGP LU Prefix SID redistribution between IGP domains (MX480, MX960, PTX1000, and PTX10008)**
 - Starting in Junos OS Release 22.2R1, we support BGP LU prefix-sid redistribution between IGP domains, and installing mpls.0 stitch route to BGP-LU next-hop.

To set prefix segment attributes, use `set prefix-segment` statement at the [edit policy-options policy-statement <policy-name> term <tern-name> from] hierarchy level.

You can now enable prefix-sid redistribution between BGP and ISIS via policy configuration without specifying index under prefix-segment. To do this, use set prefix-segment redistribute option at the [edit policy-options policy-statement <policy-name> term <term-name> then] hierarchy level.

[See [prefix-segment](#).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Support for application-specific link attribute in OSPFv2 for segment routing traffic engineering (ACX753, ACX710, MX204, MX960, MX10008, and MX2020)**—Starting in Junos OS Release 22.2R1, you can advertise different te-attributes such as te-metric, delay-metric, or admin-groups for RSVP and flexible algorithms on the same link. This is done using flexible algorithm specific application-specific link attribute as defined in RFC 8920.

To configure flexible algorithm application-specific te-attribute, include the application-specific statement at the [edit protocols ospf area interface] hierarchy level and the strict-asla-based-flex-algorithm statement at the [edit protocols ospf source-packet-routing] hierarchy level.

[See [Understanding OSPF Flexible Algorithm for Segment Routing](#).]

- **BGP classful transport (CT) support for IPv6 and Segment Routing Traffic-Engineered (SR-TE) color-only support (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 22.2R1, we support BGP-CT with IPv6 and BGP service-routes with a color-only mapping community. We have also enhanced the transport-class configuration statement to provide strict resolution without falling back on best-effort tunnels.

[See [use-transport-class](#), [BGP Classful Transport \(BGP-CT\) with Underlying Colored SR-TE Tunnels Overview](#).]

- **SRv6 locator summarisation, locator anycast, and service mapping (MX Series) (MX 10008)**—Starting in Junos OS Release 22.2R1, IS-IS can summarise and advertise locator prefixes. We also support SRv6 anycast locators and service mapping. SRv6 anycast locators identify a set of topologically near nodes to forward packets addressed to an anycast address.

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

Software Installation and Upgrade

- **Secure Zero Touch Provisioning (SZTP) (MX-Series)**—Starting in Junos OS Release 22.2R1S2, you can use RFC-8572-based SZTP to bootstrap your remotely located network devices that are in a factory-default state. SZTP enables mutual authentication between the bootstrap server and the network device before the remote network device is accessed for initiating zero touch provisioning.

To enable mutual authentication, you need a unique digital voucher, which is generated based on the DevID (Digital Device ID or Cryptographic Digital Identity) of the network device. The DevID is

embedded inside the Trusted Platform Module (TPM) 2.0 chip on the network device. Juniper Networks issues a digital voucher to customers for each eligible network device.

[See [Secure ZTP Quick Start Guide](#), [Secure Zero Touch Provisioning \(SZTP\)](#) and [Generate Voucher Certificate](#).]

Subscriber Management and Services

- **BNG redundancy for DHCP subscribers using packet-triggered based recovery (MX Series)**—In Junos OS Release 22.2R1, we've introduced broadband network gateway (BNG) redundancy for DHCP subscribers using packet-triggered based recovery. This type of redundancy is also known as stateless redundancy. This feature provides simple, easy-to-use, and lightweight stateless BNG redundancy for DHCP subscriber services with minimal traffic loss.

The stateless BNG redundancy for DHCP subscribers supports dynamic C-VLAN and static VLAN models for both relay and server.

[See [BNG Redundancy for DHCP Subscribers Using Packet Triggered Based Recovery](#), [auto-configure \(IPv4\)](#), and [auto-configure \(IPv6\)](#).]

- **Define maximum number of failed attempts for a PPP service (MX Series)**—Starting in Junos OS Release 22.2R1, you can define the maximum number of failed attempts allowed while establishing the Point-to-Point Protocol (PPP) service. Configure the `max-failures` statement at the `[edit protocols ppp-service]` hierarchy level.

[See [max-failures](#).]

- **N+1 support for BNG M:N subscriber service redundancy (MX Series)**—Starting in Junos OS Release 22.2R1, a single broadband network gateway (BNG) can perform as a backup BNG for multiple primary BNGs. You can configure the backup BNG in service-activation-on-failover mode, which consumes less resources to back up the maximum number of subscribers. The service-activation-on-failover mode enables the line cards in the backup BNG to host three times more subscribers than the primary BNGs.

This enhancement significantly reduces the resources that are reserved for redundancy in the backup BNG.

[See [N+1 Support for BNG M:N Subscriber Service Redundancy](#), [redundancy \(M:N Subscriber Redundancy\)](#), and [show system subscriber-management redundancy-state interface](#).]

- **Support for guaranteed bit rate (GBR) on Junos Multi-Access User Plane (MX240, MX480, and MX960)**—Starting in Junos OS Release 22.2R1, the Junos Multi-Access User Plane has added GBR support and supports 3GPP standards for both 4G and 5G networks. The following features are added:
 - GBR support in the downlink direction and partial support in the uplink direction

- Bandwidth reservation for express and GBR traffic flows
- Mapping of transport level marking to forwarding classes
- Call admission control (CAC)
- Maximum bit rate (MBR) and GBR policers

[See [QoS in Junos Multi-Access User Plane](#).]

VPNs

- **New ARI-TS routing protocol type for IPsec VPN traffic selector routes (MX-SPC3, SRX Series firewalls, and vSRX running iked process)**—Starting in Junos OS Release 22.2R1, when an IPsec negotiation is completed using a traffic selector configuration, the routes are installed as auto route insertion for traffic selectors (ARI-TS) routes instead of static routes.

Starting in Junos OS Release 22.2R1, ARI routes are considered as a routing protocol. These routes are installed with the same route preference and metric as in the previous implementation. With this approach, you can change the default route preference of the ARI-TS routes without impacting other routing protocols. You can also change the default preference value of the ARI-TS protocol per traffic selector to override the global option.

As ARI-TS is a new protocol, you may need to update routing policy statements depending on the configuration.

- To modify the default preference value with a global scope for an ARI-TS route, use the `set protocol ipsec-traffic-selector preference pref-value` command.
- To modify the preference value at each traffic selector level—that is, to configure a local preference value for an ARI-TS route, use the `set security ipsec vpn vpn-name traffic-selector ts-name preference pref-value` command.
- To add the ARI-TS protocol as the policy option along with the existing protocols such as BGP and OSFP, use the `set policy-options policy-statement policy-name term term-name from protocol ari-ts` command.

If you've configured the preference values at both global and local levels, the local preference value takes precedence.

[See [Understanding Traffic Selectors in Route-Based VPNs](#), [ipsec-traffic-selector](#), and [traffic-selector](#).]

Additional Features

Support for the following features has been extended to these platforms.

- **BGP, OSPF, and OSPFv3 authentication and encryption using manual IPsec SA** (MX240, MX480, and MX960 with MX-SPC3, SRX Series devices and vSRX running iked process). OSPF for IPv6, also known as OSPF version 3 (OSPFv3), does not have built-in authentication to ensure that routing packets are not altered and re-sent to the router. Starting in Junos OS Release 22.2R1, you can use IPsec to encrypt and secure BGP, OSPF, and OSPFv3 packets.

To configure IPsec for BGP, OSPF, and OSPFv3, define a security association (SA) with the security-association sa-name configuration option at the [edit security ipsec] hierarchy level for both MX Series and SRX Series platforms. You then apply the configured SA to the BGP, OSPF, and OSPFv3 configurations.

[See [security-association](#).]

To view the configured IPsec SAs for BGP, OSPF, and OSPFv3:

- On MX240, MX480, and MX960 with MX-SPC3, and on SRX Series devices and vSRX running the iked process, use the `show security ipsec control-plane-security-associations` command.

[See [show security ipsec control-plane-security-associations](#).]

- On MX240, MX480, and MX960 routers with MS-MPC/MS-MIC, use the `show ipsec security-associations` command.

[See [show ipsec security-associations](#).]

- On SRX Series devices running the kmd process, use the `show security ipsec security-associations` command.

[See [show security ipsec security-associations](#).]



NOTE: We do not support this feature with BGP, OSPF, and OSPFv3 over the secure tunnel (st0) interface.

[See [Understanding OSPFv3 Authentication, Using IPsec to Secure OSPFv3 Networks \(CLI Procedure\)](#), and [Example: Configuring IPsec Authentication for an OSPF Interface](#).]

- **Collect ON_CHANGE BGP RIB telemetry statistics and BGP neighbor telemetry with sharding** (MX Series, PTX Series and QFX Series)

[See [Telemetry Sensor Explorer](#).]

- **Lightweight PE-CE Loop Detection on EVPN-VXLAN Fabrics** (EX4400-48MP, EX4400-48P, EX9200, MX240, MX480, MX960, and MX10003)

[See [EVPN-VXLAN Lightweight Leaf to Server Loop Detection](#).]

- **Layer 2 Protocol Tunneling (L2PT)** (MX240, MX480, and MX960 with MPC10E-15C-MRATE and MPC10E-10C-MRATE; MX2010 and MX2020 with MX2K-MPC11E)

[See [Layer 2 Protocol Tunneling](#).]

- **Support for EVPN-VPWS (MX240, MX480, MX960, MX2010, and MX2020 with MPC10E line card)**
 - We've extended support for EVPN-VPWS to the listed platforms as follows:
 - EVPN-VPWS with single-active or all-active multihoming capabilities and inter-autonomous system (AS) options associated with BGP-signaled VPNs.
 - EVPN VPWS with Pseudowire Headend Termination (PWHT) on Layer 3 VPN with single-active or all-active multihoming.
 - EVPN VPWS with PWHT on VPLS with single-active multihoming.

EVPN VPWS with flexible cross connect (FXC) is not supported on the listed platforms in Junos OS Release 22.2R1.

[See [Overview of VPWS with EVPN Signaling Mechanisms](#) and [Overview of Headend Termination for EVPN VPWS](#).]

- **Support for flexible tunnel interfaces** (MX304, MX10008, and MX10016)

[See [Flexible Tunnel Interfaces Overview](#).]

- **Support for Routing-Engine-based traffic sampling** (MX10K-LC9600 line card)

[See [Configuring Traffic Sampling on MX, M and T Series Routers](#).]

- **Supported transceivers, optical interfaces, and DAC cables** Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.
- **Symmetric integrated routing and bridging (IRB) with EVPN Type 2 routes** (EX4400, EX4650, EX9204, EX9208, EX9214, MX Series, vMX, QFX5110, QFX5120, QFX10002, QFX10002-60C, QFX10008, and QFX10016). We support this feature only with MAC-VRF EVPN routing instance configurations and MAC-VRF service types `vlan-based` and `vlan-aware`. [See [Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics](#) and [irb-symmetric-routing](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control](#) | 94

- General Routing | [94](#)
- Interfaces and Chassis | [96](#)
- Layer 2 Ethernet Services | [96](#)
- MPLS | [97](#)
- Network Management and Monitoring | [97](#)
- Routing Protocols | [98](#)
- User Interface and Configuration | [99](#)
- VPNs | [100](#)

Learn about what changed in this release for MX Series.

Authentication and Access Control

- **SHA-1 password format deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX and vSRX)**—We've removed the sha1 option at the [edit system login password format] hierarchy level because SHA-1 is no longer supported for plain-text password encryption.

General Routing

- **Modified show ancp subscriber details output fields (MX Series)**—As the access loop encapsulation is transport independent it can be either passive optical network (PON) or DSL TLV. Hence, the show ancp subscriber details output field should not tag the details as a DSL TLV. Therefore, we've modified the existing DSL Line Data Link, DSL Line Encapsulation, and DSL Line Encapsulation Payload output fields to the following respectively:
 - Access Loop Encapsulation Data Link
 - Access Loop Encapsulation Encapsulation1
 - Access Loop Encapsulation Encapsulation2

See [[show ancp subscriber](#).]

- **Router advertisement module status on backup Routing Engine (MX Series)**—The router advertisement module does not function in the backup Routing Engine as the Routing Engine does

not send an acknowledgment message after receiving the packets. Starting in this Junos OS Release, you can view the router advertisement module information using the `show ipv6 router-advertisement` operational command.

See [[show ipv6 router-advertisement](#).]

- **Support for DDoS protocol (MX10008)**—We've enabled the DDoS protocol support at the `edit system ddos-protection` hierarchy level for MX10008 devices. In earlier releases, the MX10008 devices did not support these DDoS protocol statements.
 - Filter-action
 - Virtual-chassis
 - Ttl
 - Redirect
 - Re-services
 - Re-services-v6
 - Rejectv6
 - L2pt
 - Syslog
 - Vxlan

See [[protocols \(DDoS\)](#).]

- **No support for PKI operational mode commands on the Junos Limited version (MX Series routers, PTX Series routers, and SRX Series devices)**—We do not support `request`, `show`, and `clear` PKI-related operational commands on the limited encryption Junos image ("Junos Limited"). If you try to execute PKI operational commands on a limited encryption Junos image, then an appropriate error message is displayed. The `pkid` process does not run on Junos Limited version image. Hence, the limited version does not support any PKI-related operation.
- **The `<request-system-zeroize/>` RPC response indicates when the device successfully initiates the requested operation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the `<request-system-zeroize/>` RPC successfully initiates the zeroize operation, the device emits the `<system-zeroize-status>zeroizing re0</system-zeroize-status>` response tag to indicate that the process has started. If the device fails to initiate the zeroize operation, the device does not emit the `<system-zeroize-status>` response tag.
- OpenConfig container names for Point-to-Multipoint per interface ingress and egress sensors are modified for consistency from "signalling" to "signaling".

- For Access Gateway Function (AGF) statistics, consistency changes are implemented for specific leaf values in telemetry data to match field values in Junos CLI operational mode commands.

AGF NG Application Protocol (NGAP) data streamed to a collector and viewable from the Junos CLI now displays "ngap-amf-stats-init-ctx-setup-failure" and Access and Mobility Function (AMF) overload state now displays "On, Off".

- Instance type change is not permitted from default to L3VRF in open configuration (ACX Series, EX Series, MX Series, QFX Series, SRX Series, vMX, and vSRX)**—DEFAULT_INSTANCE is the primary instance that runs when there is no specific instance type configured in the route set routing-options. Any instance you explicitly configure is translated into set routing-instance r1 routing-options. The issue appears in translation, when you change instance type DEFAULT_INSTANCE (any instance to DEFAULT_INSTANCE) to L3VRF or L3VRF to DEFAULT_INSTANCE. As a result, such changes are not permitted. Additionally, DEFAULT_INSTANCE can only be named DEFAULT, and DEFAULT is reserved for DEFAULT_INSTANCE, therefore allowing no such changes.
- The request vmhost jdm login option visible to non-root users in in-chassis Junos node slicing (MX2010, MX2020, MX480, MX960, MX2008)**—The login option under the request vmhost jdm CLI is visible to non-root users. This option was earlier visible only to users with the root privileges. Though this option is now visible to all users, only root users can log in to JDM. If a non-root user attempts to log in, the software displays the following warning message:

warning: Login as ?root? to use this functionality

[See [request vmhost jdm login \(In-Chassis Model\)](#).]

Interfaces and Chassis

- Display the donor details of the IPv6 borrower interface? The output for the show interfaces command now displays the donor details of the IPv6 borrower interface.

[See [show interfaces](#).]

Layer 2 Ethernet Services

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

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

[See [show system subscriber-management statistics](#).]

MPLS

- Starting with Junos OS 16.1 the MPLS EXP bits transmitted in self ping messages are set based on the DSCP/ToS setting of the corresponding IP packet.
- When defining a constrained path LSP using more than one strict hop belonging to the egress node, the first strict hop must be set to match the IP address assigned to the egress node on the interface that receives the RSVP Path message. If the incoming RSVP Path message arrives on an interface with a different IP address the LSP is rejected.
- **Disable sending of RSVP hellos over a bypass LSP (MX Series)**—Junos routers send RSVP hello packets over a bypass LSP (when one is present), instead of the IGP next hop. To return to the original behavior specify the no-node-hello-on-bypass option.

[See [no-node-hello-on-bypass](#).]

Network Management and Monitoring

- **DES deprecation for SNMPv3**—The Data Encryption Standard (DES) privacy protocol for SNMPv3 is deprecated due to weak security and vulnerability to cryptographic attacks. For enhanced security, configure the triple Data Encryption Standard (3DES) or the Advanced Encryption Standard (CFB128-AES-128 Privacy Protocol) as the encryption algorithm for SNMPv3 users.

[See [privacy-3des](#) and [privacy-aes128](#).]

- **Changes when deactivating or deleting instances of the ephemeral configuration database (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The following changes apply when you deactivate or delete ephemeral database instances in the static configuration database:
 - When you deactivate the entire [edit system configuration-database ephemeral] hierarchy level, the device deletes the files and corresponding configuration data for all user-defined ephemeral

instances. In earlier releases, the files and configuration data are preserved; however, the configuration data is not merged with the static configuration database.

- When you delete an ephemeral instance in the static configuration database, the instance's configuration files are also deleted. In earlier releases, the configuration files are preserved.
- You can delete the files and corresponding configuration data for the default ephemeral database instance by configuring the `delete-ephemeral-default` statement in conjunction with the `ignore-ephemeral-default` statement at the `[edit system configuration-database ephemeral]` hierarchy level.

[See [Enable and Configure Instances of the Ephemeral Configuration Database](#).]

- **Limits increased for the `max-datasize` statement (ACX Series, PTX Series, and QFX Series)**—The `max-datasize` statement's minimum configurable value is increased from 23,068,672 bytes (22 MB) to 268,435,456 bytes (256 MB), and the maximum configurable value is increased from 1,073,741,824 (1 GB) to 2,147,483,648 (2 GB) for all script types. Furthermore, if you do not configure the `max-datasize` statement for a given script type, the default maximum memory allocated to the data segment portion of a script is increased to 1024 MB. Higher limits ensure that the device allocates a sufficient amount of memory to run the affected scripts.

[See [max-datasize](#).]

- **Changes to the NETCONF `<edit-config>` RPC response (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the `<edit-config>` operation returns an error, the NETCONF server does not emit a `<load-error-count>` element in the RPC response. In earlier releases, the `<edit-config>` RPC response includes the `<load-error-count>` element when the operation fails.
- **Change in unnumbered-address support for GRE tunnel**—Starting in Junos OS Release 24.4R1, there is a behavioural change in unnumbered-address support for GRE tunnel with IPV6 family and display donor interface for both IPV4 and IPV6 families of GRE tunnel. You can view interface donor details under `show interfaces` hierarchy level.

[See [show interfaces](#).]

Routing Protocols

- **The `RPD_OSPF_LDP_SYNC` message not logged**—On all Junos OS and Junos OS Evolved devices, when an LDP session goes down there is a loss of synchronization between LDP and OSPF. After the loss of synchronization, when an interface has been in the holddown state for more than three minutes, the system log message with a warning level is sent. This message appears in both the messages file and the trace file.

However, the system log message does not get logged if you explicitly configure the hold-time for ldp-synchronization at the [edit protocols ospf area area id interface interface name] hierarchy level less than three minutes. The message is printed after three minutes.

- To achieve consistency among resource paths, the resource path `/mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter[ip-addr='address']/state/counters[name='name']/out-pkts/` is changed to `/mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter[ip-addr='address']/state/counters[name='name']/`. The leaf "out-pkts" is removed from the end of the path, and "signalling" is changed to "signaling" (with one "l").
- When the `krt-nexthop-ack` statement is configured, the RPD will wait for the next hop to get acknowledged by PFE before using it for a route. Currently, only BGP-labeled routes and RSVP routes support this statement. All other routes will ignore this statement.
- **SSH TCP forwarding disabled by default**—We've disabled the SSH TCP forwarding feature by default to enhance security. To enable the SSH TCP forwarding feature, you can configure the `allow-tcp-forwarding` statement at the [edit system services ssh] hierarchy level.

In addition, we have deprecated the `tcp-forwarding` and `no-tcp-forwarding` statements at the [edit system services ssh] hierarchy level.

[See [services \(System Services\)](#).]

User Interface and Configuration

- **Load JSON configuration data with unordered list entries (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos schema requires that list keys precede any other siblings within a list entry and appear in the order specified by the schema. Junos devices provide two options to load JSON configuration data that contains unordered list entries:
 - Use the `request system convert-json-configuration` operational mode command to produce JSON configuration data with ordered list entries before loading the data on the device.
 - Configure the `reorder-list-keys` statement at the [edit system configuration input format json] hierarchy level. After you configure the statement, you can load JSON configuration data with unordered list entries, and the device reorders the list keys as required by the Junos schema during the load operation.
 - When you configure the `reorder-list-keys` statement, the load operation can take significantly longer to parse the configuration, depending on the size of the configuration and number of lists. Therefore, for large configurations or configurations with many lists, we recommend using the `request system convert-json-configuration` command instead of the `reorder-list-keys` statement.

[See [json](#) and [request system convert-json-configuration](#)]

- **Junos XML protocol Perl modules deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—We no longer provide the Junos XML protocol Perl client for download. To use Perl to manage Junos devices, use the NETCONF Perl library instead.

[See [Understanding the NETCONF Perl Client and Sample Scripts..](#)]

- **Persistent CLI timestamps**—To have a persistent CLI timestamp for the user currently logged in, enable the `<codeph>set cli timestamp</codeph>` operational command. This ensures the timestamp shows persistently for each new line of each SSH session for the user or class until the configuration is removed. To enable timestamp for a particular class with permissions and format for different users, configure the following statements:

```
set system login class <variable>class name</variable> permissions <variable>permissions</variable> set system
login class <variable>class name</variable> cli timestamp set system login user username class <variable>class
name</variable> authentication plain-text-password
```



NOTE: The default timestamp format is `%b %d %T`. You can modify the format per your requirements. For example, you can configure the following statement:

To enable timestamp for a particular user with default class permissions and format, configure the following statements:
`set system login user username class <variable>class name</variable> authentication plain-text-password set system login user <variable>username</variable> cli timestamp`

VPNs

- **Changes to `show mvpn c-multicast` and `show mvpn instance outputs`**—The FwdNh output field displays the multicast tunnel (mt) interface in the case of Protocol Independent Multicast (PIM) tunnels.

[See [show mvpn c-multicast](#).]

Known Limitations

Learn about known limitations in Junos OS Release 22.2R1 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

- On all MX devices, traffic might be interrupted when you change the configuration from AMS warm-standby to AMS deterministic NAT. [PR1597386](#)

MPLS

- With local reversion ON, there is a possibility of transit router not informing headend of RSVP disabled link when link flaps more than once. As a workaround, remove the local-reversion configuration. [PR1576979](#)
- The automatic sorting of configuration entries does not work if defined under a group. [PR1637730](#)

Platform and Infrastructure

- Major alarm with **XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8)** logs appears. [PR1303489](#)
- On MX series devices, under Ethernet VPN (EVPN) environment, packets routed using IRB interface could not be fragmented due to media maximum transmission unit (MTU) problem. [PR1522896](#)
- Deactivating services rpm/rpm-tracking does not remove the tracked route from the routing or forwarding tables. [PR1597190](#)
- Routing Engine-based BFD sessions might flap during switchover when there are large number of BFD, IS-IS, OSPF and LDP packets to be sent out. [PR1600684](#)

Routing Policy and Firewall Filters

- When a hierarchy specified in apply-path is configured in an ephemeral instance, then rpd might not be able to update the prefix-list for the configuration in ephemeral instance. [PR1636390](#)

Routing Protocols

- When we have high scale, the openconfig telemetry sensor /bgp-rib/ used in periodic streaming will cause high cpu usage by RPD. [PR1625396](#)
- When routing-options transport-class fallback is not configured, do not configure more than 10 transport-classes or advertise more than 10 distinct colors in SR-TE or FlexAlgo. [PR1648490](#)

Open Issues

Learn about open issues in Junos OS Release 22.2R1 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

- If a vmhost snapshot is taken on an alternate disk and no further vmhost software image is upgraded, the expectation is that if the current vmhost image gets corrupted, the system boots with the alternate disk so the user can recover the primary disk to restore the state. However, the host root file system and the node boots with the previous vmhost software instead of the alternate disk. [PR1281554](#)
- On MX Series devices with MPC7E, MPC8E, or MPC9E installed, if optics QSFP-4X10GE-LR from vendor (subset of modules with part number 740-054050) is used, the link might flap. [PR1436275](#)
- VXLAN VNI (multicast learning) scaling on QFX5110 traffic issue is seen from VXLAN tunnel to Layer 2 interface. [PR1462548](#)
- With NAT/Stateful-firewall/TCP tickle (enable by default) configured on MS-MPC/MS-MIC, the vmcore process crashes sometimes along with mspmand crash might happen if large-scale traffic flows (e.g. million flows) are processed by it. [PR1482400](#)
- When there are HW link errors occurred on all 32 links on an FPC 11. Because of these link errors, all FPCs reported destination errors towards FPC 11 and FPC 11 was taken offline with reason **offline due to unreachable destinations**. [PR1483529](#)
- The WAN-PHY interface continuously flaps with the default hold-time down of value 0. This is not applicable to an interface with the default framing LAN-PHY. [PR1508794](#)
- When an AMS physical interface is configured for the first time or any member of the AMS bundle is removed or added, the PICs on which the members of AMS bundle are present go for a reboot. There is a timer running in the AMS kernel which is used as a delay for the PIC reboot to complete and once that timer expires, AMS assumes that the PICs might have been rebooted, and it moves into next step of AMS finite state machine (FSM). In scaled scenarios, this rebooting of the PIC is delayed due to DCD. This is because when a PIC goes down, DCD is supposed to delete the physical interfaces on that PIC and the PIC reboot happens. But DCD is busy processing the scaled configuration and the physical interface deletion is delayed. This delay is much greater than the timer running in AMS kernel. When the timer expires, the FSM in AMS kernel incorrectly assumes the PIC reboot would be completed by then, but the reboot is still pending. By the time DCD deletes this physical interface, the AMS bundles are already up. Because of this, there is a momentary flap of the bundles. [PR1521929](#)

- Due to BRCM KBP issue route lookup might fail. Need to upgrade KBP to address this issue. [PR1533513](#)
- The riot might crash due to a rare issue if vMX run in the performance mode. [PR1534145](#)
- FPC might generate a core file if flap-trap-monitor feature under set protocols oam ethernet cfm performance-monitoring sla-iterator-profiles is used and performance monitoring flap occurs. [PR1536417](#)
- In scaled MX2020 devices, with vrf localisation enabled, 4 million nexthop scale, 800,000 route scale. FPCs might go offline on GRES. Post GRES, router continues to report many fabric related CM_ALARMS. FPC might continue to reboot and not come online. Rebooting the primary and backup Routing Engine will help recover and the router gets stable. [PR1539305](#)
- The mspmand process leaks memory in relation to the MX Series telemetry reporting the following error message: **RLIMIT_DATA exceed**. [PR1540538](#)
- 5M DAC connected between QFX10002-60C and MX2010 devices doesn't link up. But with 1M and 3M DAC, this interoperation works as expected. Also it is to be noted on QFX10002-60C and ACX Series devices or traffic generator, the same 5M DAC works seamlessly. There seems to be a certain SI or link-level configuration on both QFX10002-60C and MX2010 devices which needs to be debugged with the help from HW and SI teams and resolved. [PR1555955](#)
- The SyncE to PTP transient response is a stringent mask to be met with two way time error. The SyncE to PTP transient response mask might not be met for MPC7E-1G and MPC7E-10G line cards. [PR1557999](#)
- Support switchover on routing-crash configuration statement during abnormal termination of rpd. [PR1561059](#)
- Due to a race condition, the show multicast route extensive instance instance-name command output can display the session status as invalid. Such an output is a cosmetic defect and not an indicative of a functional issue. [PR1562387](#)
- Interface hold time needs to be configured to avoid the additional interface flap. [PR1562857](#)
- Copying files to /tmp/ causes a huge JTAKSCHED_SLIP. Copy files to /var/tmp/ instead. [PR1571214](#)
- This issue is caused by /8 pool with block size as 1. When the configuration is committed, the block creation utilizes more memory causing NAT pool memory shortage, which is currently being notified to the customer with syslog tagged RT_NAT_POOL_MEMORY_SHORTAGE. [PR1579627](#)
- In a fully loaded devices, at times, firewall programming was failing due to scaled prefix configuration with more than 64800 entries. However, this issue is not observed in development setup. [PR1581767](#)

- On all devices running Junos OS Release 19.1R3-S5-J3, the subscriber IFL(logical interface) might be in a stuck state after the ESSM (Extensible Subscriber Services Manager) deletion. [PR1591603](#)
- Pim VXLAN does not work on the TD3 chipsets that enables the VXLAN flexflow. [PR1597276](#)
- On MX2010 and MX2020 Series devices: MPC11E: Unified ISSU is not supported for software upgrades from 21.2 to 21.3 and 21.4 releases due to a flag day change.[PR1597728](#)
- On the MX10008 and MX10016 devices, during Routing Engine switchover, if there is a burst of ICMP, BFD, SSH, FTP, TELNET, and RSVP packets (~18,000 pps), then the new backup Routing Engine might restart.[PR1604299](#)
- On aggregate Ethernet interfaces with some of the member links part of MPC10 or MPC11, and other member links part of other MPC type (MPC1 up to MPC9), if you delete an "ae" interface, other "ae" interfaces may experience unicast packet loss.[PR1604450](#)
- On MX-VC (Virtual Chassis) platforms with MS-MPC or SPC3 service cards and AMS (Aggregated Multi-Service), traffic on the line card in the backup chassis might not be load-balanced properly due to timing conditions. This works well on the line card in the primary chassis. There might be traffic loss when interfaces are not properly balanced.[PR1605284](#)
- NPU sensor path for subscription is: /junos/system/linecard/npu/memory/ It's output would contain info like: system_id:wf-mt-ranier component_id:4 path:sensor_1004_1_1:/junos/system/linecard/npu/memory/:/junos/system/linec ard/npu/memory/:aftd-trio sequence_number:1 timestamp:1639179017148 .. kv { key:property[name='mem-util-firewall-fw-bytes-allocated']/state/value int_value:9064 } kv { key:property[name='mem-util-firewall-fw-allocation-count']/state/value int_value:94 } kv { key:property[name='mem-util-firewall-fw-free-count']/state/value int_value:0 } kv { key:property[name='mem-util-firewall-inline-jflow-sample-rr-(dfw)-bytes-all ocated']/state/value int_value:131160 } kv { key:property[name='mem-util-firewall-inline-jflow-sample-rr-(dfw)-allocatio n-count']/state/value int_value:6 } kv { key:property[name='mem-util-firewall-inline-jflow-sample-rr-(dfw)-free-coun t']/state/value int_value:0 } kv { key:property[name='mem-util-firewall-inline-jflow-sample-nh-(dfw)-bytes-all ocated']/state/value int_value:16 } kv { key:property[name='mem-util-firewall-inline-jflow-sample-nh-(dfw)-allocatio n-count']/state/value int_value:1 } kv { key:property[name='mem-util-firewall-inline-jflow-sample-nh-(dfw)-free-coun t']/state/value int_value:0 } kv { key:property[name='mem-util-firewall-fw-strided-bytes-allocated']/state/val ue int_value:9064 } kv { key:property[name='mem-util-firewall-fw-strided-allocation-count']/state/val ue int_value:94 } kv { key:property[name='mem-util-firewall-fw-strided-free-count']/state/value int_value:0 } kv { key:property[name='mem-util-counters-fw-counter-bytes-allocated']/state/val ue int_value:16416 } kv { key:property[name='mem-util-counters-fw-counter-allocation-count']/state/val ue int_value:3 } .. The (VTY) CLI output is: root@wf-mt-ranier-fpc4:pfe> show npu memory info | match firewall mem-util-firewall-ro-edmem-size 20971520 mem-util-firewall-ro-edmem-allocated 294912 mem-util-firewall-ro-edmem-utilization 1 mem-util-firewall-ro-edmem-size 20971520 mem-util-firewall-ro-edmem-allocated 294912 mem-util-firewall-ro-edmem-utilization 1 mem-util-firewall-ro-edmem-size 20971520 mem-util-firewall-ro-edmem-allocated 294912 mem-util-firewall-ro-edmem-utilization 1 mem-util-firewall-ro-edmem-size 20971520 mem-util-firewall-ro-edmem-allocated 294912 mem-util-firewall-ro-edmem-utilization 1. [PR1606791](#)

- On all MX devices, in a subscriber management environment, new subscribers might not connect if CoS (Class of service) CR-features (Classifier Rewrite) are used by the VBF (Variable Based Flow) service. The reference count mismatching between RE (Routing Engine) and VBF is caused by VBF flow VAR CHANGE failure. [PR1607056](#)
- When rpd sends INH deletion or additions out of order (rarely occurs) message to backup rpd, the rpd crashes and generates a core file. [PR1607553](#)
- Duplicate syslog messages gets displayed for IPv4 and IPv6 sessions after the Configure NAT Services with 2 service sets (next-hop style) one for NAPT44 and another for NAPT64. [PR1614358](#)
- The errors are displayed with following next-hop hieINH->COMPNH->UCAST->AE_IFL. During AE-IFL flaps control detects and initiate MBB. Its possible by that Packet Forwarding Engine can see an compNH->ucast with ae-ifl down resulting into these error messages but this is only transient. There is no functionality impact due to these error messages. [PR1617388](#)
- On MX480 devices, the ntf-agent services are not running and TCP connection is refused between router and ipfix-collector. [PR1626505](#)
- Tunnel interface statistics displays incorrect values when jflow sampling is enabled. [PR1627713](#)
- For MX204 and MX2008 "VM Host-based" devices, starting with Junos OS 21.4R1 Release or later, ssh and root login is required for copying line card image (chspmb.elf for MX2008) from Junos VM to Linux host during installation. The ssh and root login are required during installation. Use **deny-password** instead of **deny** as default root-login option under ssh configuration to allow internal trusted communication. Ref <https://kb.juniper.net/TSB18224> [PR1629943](#)
- The fabric statistics counters are not displayed in the output of **show snmp mib walk ascii jnxFabricMib**. [PR1634372](#)
- On Junos OS platforms, high BGP scale with flapping route and BGP Monitoring Protocol (BMP) collector or station is very slow might cause rpd crash due to memory pressure. [PR1635143](#)
- The USB device on MX304 device can be accessed from host linux instead of Junos OS (as is usually done on most other platforms). MX304 device is similar to PTX1000 device in this respect.
- On MPC10E cards, upon many very quick link down and up events in msec range might not always able to drain all traffic in the queue. This causes lost of traffic going through the interface. Traffic volume and class-of-service configuration does influence the exposure. [PR1642584](#)
- WIth PTPolIPv6 on MPC2E 3D EQ, PTP slave stays in acquiring state. [PR1642890](#)
- Class-of-service buffer-size exact config is not supported. The respective configured queue will still use the shared-pool. [PR1644355](#)
- Committing configuration changes during the PFE (Packet Forwarding Engine) reset pause window (when PFE is disabled, yet the PFE reset proper has not started yet) has the potential of causing

errors and traffic loss. In particular, configuration changes that result in re-allocating policers (which are HMC-based) might lead to traffic being entirely policed out (i.e. not flowing). Once the PFE reset procedure has started config changes ought to be avoided until the procedure is completely done.[PR1644661](#)

- Run with BB device enabled using CLI command in configuration for IPoE and PPPoE access models. [PR1645075](#)
- Configuring MPC11 in 4x100G and keeping peer in 400G mode, link comes up on peer while staying down on local end. [PR1653946](#)
- When interop with the following systems, flow control must be enabled when MACsec is configured on the peer system. Because on these systems, flow control is forced to be on regardless of the CLI provisioning. [PR1655712](#)
- Core seen intermittently where random grpc stack crash is observed. License service will need to be restarted.[PR1656975](#)
- Node-index in link key is a short and cannot hold when the to node's index is more than 32 bits long. Once this index exceeds its limitation, SR-TE LSP will become down due to **Compute Result failure**.[PR1657176](#)
- TOS(DSCP+ECN) bits does not get copied from the inner Layer 3 header to outer VXLAN header at the Ingress VTEP. Because of this in the core, ECN marking and DSCP classification does not work.[PR1658142](#)
- On GNF, no streaming data received for /telemetry-system/subscriptions/dynamic-subscriptions/.[PR1661106](#)
- Few ARP entries are not resolved for IRB interface IP when IRB is configured under VPLS routing instances.[PR1662882](#)
- MX10008 with MX10000-LC2101 Linecard(s) supports *PTP* only with JNP10008-SF Switch Fabric Board(s), *PTP* currently does not work with JNP10008-SF2 Switch Fabric Board(s).[PR1664569](#)
- On all Junos OS platforms, link-degrade functionality needs to be supported and manually configured. Link degrade is manually configured to monitor for link error or issues. Once the error is observed, the link goes down. [PR1664978](#)
- RE0 to RE1 interface EM4 MTU is changed to 9192 bytes. If one of the Routing Engines does not have this fix, Routing Engine synchronisation fails. Due to this reason, ISSU will not work. In such scenario, cold image upgrade should be done.[PR1665690](#)
- After configuring the warm-standby option, you must wait for three minutes before Routing Engine switchover. [PR1623601](#)

- In case of routing instance type EVPN or EVPN-VPWS, the system automatically creates one default routing instance apart from EVPN and/or EVPN-VPWS routing instance. In the output of the show snmp mib walk jnxVpnInfo command, the number of configured routing instances are always one more than the number of EVPN and/or EVPN-VPWS instances configured in the system. [PR1659466](#)
- If the interface configuration for fxp0 and lo0 gets deleted and you commit the configuration, the configuration of the internal network interfaces to spmb process also gets deleted. This results in SPMB process going down and generating major alarms. Ensure that the configuration should always have fxp0 or lo0 available. You should not commit any configuration that does not have interfaces configuration for fxp0 or lo0. [PR1640746](#)

EVPN

- In a PBB-EVPN environment, the ARP suppression feature, which is not supported by the PBB might be enabled unexpectedly. This might cause MAC addresses of remote CEs not to be learned and hence traffic loss might be seen. [PR1529940](#)
- This is a case where interface is disabled and comes up as CE after a timeout. A manual intervention of clear CE interface command should restore this. This can be a workaround: 1) clear auto-evpn ce-interface <interface-name> 2) configure edit activate <interface-name> family inet inet6. We can fix this in phase 2 by keeping some persistent state on a interface being a core facing interface in some incarnation. [PR1630627](#)

Flow-based and Packet-based Processing

- When customer perform unified ISSU with security VRF-group configuration, the unified ISSU cannot be completed successfully. [PR1661935](#)

Forwarding and Sampling

- When GRES is triggered by SSD hardware failure, the syslog error of **rpd[2191]: krt_flow_dfwd_open,8073: Failed connecting to DFWD, error checking reply - Operation timed out** might be seen. Issue can be recovered by restarting the dfwd daemon. [PR1397171](#)
- The **fast-lookup-filter** with match not supported in FLT Hardware might cause the traffic drop. [PR1573350](#)

Layer 2 Features

- Adding one more sub-interface logical interface to an existing interface causes 20 to 50 milliseconds traffic drop on the existing logical interface. [PR1367488](#)

MPLS

- In MVPN case, if the nexthop index of a group is not same between primary and backup after a nsr switchover, you might see a packet loss of 250 to 400 milliseconds. [PR1561287](#)
- The ingress retries after LSP stay down for extended period of time or customer clears LSP to speed up the retry. [PR1631774](#)

Network Management and Monitoring

- When maximum-password-length is configured and user tries to configure password whose length exceeds configured maximum-password-length, error is thrown, along with error **ok** tag is also emitted. (Ideally **ok** tag should not be emitted in an error scenario.) The configuration does not get committed. [PR1585855](#)
- A minor memory leak is seen in the event-daemon process when multiple GRES switchovers are performed. [PR1602536](#)
- The mgd process might crash when you configure an invalid value for identityref type leafs or leaf-lists while configuring Openconfig or any other third-party YANG. The issue occurs with JSON and XML loads. [PR1615773](#)

Platform and Infrastructure

- DRouting Engine switchover interface flap might be seen along with scheduler slippage. [PR1541772](#)
- If you use the source-address NTP configuration parameter and issue the command `set ntp date` from the CLI, packets are sent with the source address of the outgoing interface rather than the manually configured IP address. Typically the manually configured IP address would be a loopback address. The problem does not apply to automatically generated NTP poll packets. [PR1545022](#)
- TWAMP-Light is supported on MX Series devices. CLI configuration support will be disabled on all other platforms. Do not use the control-type light under platforms where this feature is not supported. Currently, IPv4 and IPv6 twamp-light is supported on the platforms using TRIO and PE chipsets. [PR1603128](#)
- Using static LSP(labeled switched path) configuration, the child node is not removed from the flood composite when the core interface goes down. [PR1631217](#)
- With given multi dimensional scale, if configuration is removed and restored continuously for more than 24 times, MX Trio based FPC might crash and restart. During the reboot, there can be traffic impact if backup paths are not configured. [PR1636758](#)

Routing Protocols

- On MX devices, initial multicast register packets might get dropped, this might affect multicast services. [PR1621358](#)
- When filter is configured through open configuration and bound to a routing table instance, the filter bind object is not getting published due to the absence of routing table object. Hence the filter does not work as expected since the traffic does not hit the filter. [PR1644421](#)
- RFC 8950/RFC 5549, permits the advertisement of a BGP Nexthop of a different family (e.g. IPv6) than the NLRI address family (e.g. IPv4). The mapping of possible address families that can be used are exchanged using BGP Capabilities. The BGP Capabilities specification, RFC 5492, recommends that a single capability TLV of a given type is advertised when multiple elements within that TLV are present. That RFC also permits multiple capabilities of the same type to be advertised for multiple elements for backward compatibility. Junos BGP handling of the BGP extended nexthop capability did not handle multiple capabilities of the same code point when multiple extended nexthop capabilities were present. It incorrectly kept only the last one sent. [PR1649332](#)
- Device having three routing-instance with matching IMPORT & EXPORT RT policy and when we configured auto-export in two VRF then routes from third VRF (Auto-export not configured) The route is leaking incorrectly into other two VRF with auto-export. [PR1665094](#)

VPNs

- In some scenario (for example, configuring firewall filter), routers might show obsolete IPsec SA and NHTB entry even when the peer tear down the tunnel. [PR1432925](#)
- Tunnel debugging configuration is not synchronized to the backup node. It needs to be configured again after RG0 failover. [PR1450393](#)
- When using Group VPN, in certain cases, the PUSH ACK message from the group member to the group key server might be lost. The group member can still send rekey requests for the TEK SAs before the hard lifetime expiry. Only if the key server sends any new PUSH messages to the group members, those updates would not be received by the group member since the key server would have removed the member from registered members list. [PR1608290](#)

Resolved Issues

Learn about the issues fixed in this release for MX Series.

Class of Service (CoS)

- The cosd process might not be able to send unbinds for rewrites post a certain sequence of operations are performed. [PR1649510](#)
- Interface burst size becomes low in Packet Forwarding Engine, when rate-limit-burst configuration statement is removed. [PR1650089](#)
- Hierarchical class of service (HCOS) might not work for LT interfaces configured on PIC 2 and PIC 3 of MPC5E/MPC6E. [PR1651182](#)
- Creating aggregate Engine interfaces in per-unit-scheduler mode and committing COS config on aggregate Engine IFLs in a single commit can lead to race-conditions. [PR1656441](#)

EVPN

- Few ARP/ND/MAC entries for VLANs are missing with MAC-VRF configuration. [PR1609322](#)
- The rpd might crash when moving an interface from VPLS to EVPN-VPWS instance. [PR1632364](#)
- IRB might not send out arp-reply if no-arp-suppression is configured. [PR1646010](#)
- The DF and BDF both might be Up/Forwarding in EVPN Multihoming single-active scenarios. [PR1647734](#)
- The spine might have stale vtep entry for the ESI even though the host MAC is not advertised by the leaf. [PR1648368](#)
- EVPN VXLAN Type 5 does not work with asymmetric VNI configuration. [PR1652339](#)

Forwarding and Sampling

- Delay in getting the response for clear interfaces statistics all command with scale configuration. [PR1605544](#)
- Packet loss might be reported after hitting the firewall filter on Junos OS platform. [PR1625309](#)

General Routing

- PTP packets dropped depending on multicast configuration. [PR1442055](#)
- Junos 'et-' interface stuck and remains down between two particular ports. [PR1535078](#)
- An MGD core file is generated on performing the help apropos command in configuration mode. The MGD will restart and as long as the command is not issued again. [PR1552191](#)

- The process pkid might generate core files during local certificate enrollment. [PR1573892](#)
- The CHASSISD_FRU_IPC_WRITE_ERROR: fru_send_msg: FRU GNF 2, errno 40, message too long error might appear periodically in the chassisd logs. [PR1576173](#)
- When interim logging is configured for PBA, it generates syslog messages at regular intervals. Change in the information of PBA interim syslog message, message string change from **allocates port block** to **interim port block**. [PR1582394](#)
- NAT EIM mapping gets created in the FTP ALG child sessions. [PR1587849](#)
- USP-SPC3: SESSION CLOSE Termination reason is "other" (very generic reason) when the server initiated sessions are closed due to PCP Life time expiry. [PR1588785](#)
- Implement FW reload option to Alfaromeo LC in CLI. [PR1594579](#)
- The mspmand daemon memory leak might be observed after the HA primary goes down. [PR1598356](#)
- The dcpfe process generates core files while testing ISSU from Junos OS Release 21.1R1.11 to Junos OS Release 21.2R1.7. [PR1600807](#)
- Node Slicing External JDM: After Rebooting the JDM from Shell Mode, not able to stop the JDM again. [PR1603637](#)
- On MX10008 and MX10016, Kernel error logs(tcp_timer_keep) is seen on backup Routing Engine when set system internet-options no-tcp-reset drop-all-tcp option is enabled. [PR1605255](#)
- VVM host platforms might boot exactly 30 minutes after executing request vmhost halt command. [PR1605971](#)
- **WO-0: OGE0 dequeue watermark hit** might seen with Layer 2 related configuration and receiving jumbo-frame packets. [PR1606967](#)
- IPv6 link local BFD session might not come up if there is no child link of an aggregated Ethernet mapped to pfe inst 0. This issue is applicable to MPC9 and below MX Series-based line cards. [PR1607077](#)
- MPC6E 3D did not comes back up after MIC offline online test. [PR1614816](#)
- Primary RE0 reloaded unexpectedly and new primary RE1 does not bring up IS-IS or LDP adjacencies. [PR1616114](#)
- ICMP error packet do not have relevant header when configured with DSLite and with appropriate ICMP ALG name and one UDP application name. [PR1616633](#)
- The transit IPv4-over-IPv6 encapsulated packets cannot pass through using IP over IP interface. This behavior has been seen on transit packets only. [PR1618391](#)

- Traffic might be dropped due to the TX queue memory leak on PCI interface. [PR1618913](#)
- `/interfaces/interface/subinterfaces/subinterface/state/counters` are not exported during initial synchronization for on-change. [PR1620160](#)
- FPC might crash on MX10003 when MACsec interfaces configured with bounded-delay feature are deleted in bulk. [PR1621868](#)
- [G.8275.1]: High phase jump spikes of ~4000 to 5000 ns is observed during client clock fail-over within same line card, due to configuration commit. [PR1622575](#)
- Constant increase of PCS errors might be seen on channelized port. [PR1622741](#)
- BGP Flowspec might not show counters for matching IPv6 firewall filter. [PR1623170](#)
- flowd core file is generated with TLB configuration only with the combination of MPC10 cards. [PR1624572](#)
- The mcontrol might frequently miss keepalives from backup Routing Engine. [PR1624623](#)
- Pkid crash happening due to null pointer dereferencing during local certificate verification in some cases. [PR1624844](#)
- Fabric request timeouts and fabric healing occur. [PR1625820](#)
- The primary role transfer might not be triggered on each rpd crashes if switchover-on-routing-crash is configured. [PR1625834](#)
- Traffic drop might be seen in node slicing scenario. [PR1626115](#)
- [technology/inlineservices] [core] : mx960 :: spd core is seen. [PR1626311](#)
- After configuring 4000 bridge domains messages log file is flooded with kernal messages. [PR1626381](#)
- VPLS MAC age time-out might not be applied on some MAC addresses. [PR1627416](#)
- Subscribers might face connectivity issues due to memory leak. [PR1627562](#)
- DHCPv6 server binding might not happen when LDRA is enabled along with DHCPv6 snooping. [PR1627600](#)
- Transient JSR replication errors 113 / 115 seen on disable or enable OSPF. [PR1627625](#)
- When DHCP persistence is configured with DHCP security and device reboots, the lease time values might show a high lease value post reboot. [PR1627673](#)
- On DUT with scaled MPLSVPN configuration and Junos Telemetry interface sensors configured, the stream of error messages **agentd_telemetry_uninstall_sensor: Deleting subscription from daemon**

aftsysinfo failed after mgmt_sock_retries 601, ret -1 is seen on stopping jtimon. Sensor packet drops might be seen when the error message scrolls on DUT.[PR1627752](#)

- FPC might restart with syslog filter action configured. [PR1627986](#)
- ECMP might not work properly when AMS is configured as next-hop with ECMP. [PR1628076](#)
- Invalid IP length packets encapsulated within MPLS may trigger PPE traps. [PR1628091](#)
- Tunnel-service bandwidth should not be changed when there are active subscribers. [PR1628628](#)
- DDoS filter does not classify OSPF packets as OSPF-Hello and OSPF-Data packet. [PR1628889](#)
- The pfe all setting should not be used with oc-category configurations. [PR1628964](#)
- FPC crash might be observed in the subscriber scenarios. [PR1629136](#)
- Whenever vmhost image is installed on MX10008 chassis via USB, LC9600 will eventually go offline when no interface (either loopback or management port) is configured. Configure one interface and restart chassisd process (cli> restart chassis-control) upon completion of USB installation of VMHOST image on MX10008 Routing Engines to avoid LC9600 going offline. [PR1629558](#)
- The l2ald might be stuck in **issu state** when unified ISSU is aborted. [PR1629678](#)
- The egress traffic on non-targeted ifset of subscribers might not be forwarded correctly over targeted aggregated Ethernet interface. [PR1629910](#)
- Multiple link flaps and traffic might be lost on the links. [PR1630006](#)
- LACP timeout might be observed during high CPU utilization. [PR1630201](#)
- With SCBE3+SPC3, fabric drops are seen around 10M PPS/60G TCP traffic with ~750byte packet size with IPv6 SFW on a single PIC. [PR1630223](#)
- Interfaces configured with default CTLE value using Gen 2 or more latest version of Finisar 100G SR4 optics might not come up. [PR1630300](#)
- Index of the link might get missed in the distribution table of Packet Forwarding Engines after the flap. [PR1630408](#)
- If the interface is in link up transition with Hold Up timer enable (Link down, Admin Up/ Enabled), and Packet Forwarding Engine reset occurs, the interface will come UP post Packet Forwarding Engine reset after Hold timer expiry. [PR1630793](#)
- The FPC might crash after enabling MACsec. [PR1631010](#)
- A clksync crash might be observed and PTP might get stuck. [PR1631261](#)

- PTP (Precision Time Protocol) might not lock on MX with MX-MPC2E-3D-P and MPC2E-3D LC. [PR1631274](#)
- ipv6 host route prefix match disappear from 'forwarding-table' after a ping test, ping continues to work, forwarding table entry is not shown. No impact in traffic. [PR1631607](#)
- Adverse effect on subscriber management observed after deactivating chassis pseudowire-service with active subscribers [PR1631787](#)
- Operations dependent on the SDB shared memory might be impacted. [PR1631858](#)
- [macsec] [fips MPC7E] FPC/PIC should get rebooted on fake KATs generation. [PR1632273](#)
- P2MP LSP ping and trace from bud-node might fail when the branch is on another Packet Forwarding Engine. [PR1632385](#)
- When deleting the VNI and there is another vlan-id-list with a different VNI might cause traffic loss. [PR1632444](#)
- It is noted that the single hop BFD session over aggregated Ethernet is not fully functional after exercising Packet Forwarding Engine reset feature. The BFD session was up before Packet Forwarding Engine reset operation is initiated but after the reset the BFD rx session is not fully functional. [PR1632585](#)
- The show chassis firmware does not show the revision for PIC FPGA. [PR1633187](#)
- In subscriber scenario, traffic drop might be seen when aggregate Ethernet member link is removed. [PR1634371](#)
- Traffic floods back to the source port with H-VPLS local-switching. [PR1634480](#)
- The fpc might crash on enabling port-mirroring. [PR1634570](#)
- Traffic impact might be seen when a firewall filter based policer for MPLS address family is configured on the device. [PR1634644](#)
- show network-agent statistics detail CLI output not printing expected output. after sometime [PR1634716](#)
- LACP interface might go down when a sub-interface configuration is added and committed to the aggregate Ethernet interface. [PR1634908](#)
- [MX][PFE Reset][VALE] Post Packet Forwarding Engine reset error information is not going out from show system errors active detail. [PR1635284](#)
- CFM CCM PDU is not forwarded transparently on generating a core file if the physical interface is configured under OAM protocol. [PR1635293](#)

- Some packets might be dropped inside the I2circuit when the flow label is enabled. [PR1635345](#)
- Client deadline might exceed with error after gribi route add with FIB ACK. [PR1635727](#)
- The chassisd might crash if chassis disk-partition is configured. [PR1635812](#)
- Precision Time Protocol (PTP) packets having huge correction-field (CF) value coming out from MX devices. [PR1635877](#)
- IPsec tunnel might not establish after a flap. [PR1635882](#)
- FPCs might get restarted due to either faulty PEM module. [PR1636118](#)
- CMVTS: IPsec : After core interface flap 16 out of 1000 ipsec session failed to come UP. [PR1636164](#)
- SFP-1FE-FX might not function properly on MIC-MACSEC-20G. [PR1636322](#)
- DHCP offer not getting processed in the routing instance when using LT interfaces. [PR1636579](#)
- Wrong interface statistics might be reported on MX204. [PR1636654](#)
- The interface equipped with a QSA adapter might go down. [PR1636874](#)
- FPC crash might be seen on all MX platforms with BBE subscriber. [PR1637304](#)
- Ingress PE doesn't insert Sh label for BUM traffic received on local EP ESI interface, causing packet duplication on egress PE. [PR1637703](#)
- MACsec traffic, silent packet drop might be seen when a back-to-back graceful switchover is performed. [PR1637822](#)
- Delay might be observed for the interfaces to come up after reboot or transceiver replacement. [PR1638045](#)
- BNG CUPS: ERA & OIU - Core in ouModShMemEntry during OIU modify when restarting smg-service while bouncing subscribers. [PR1638217](#)
- Packet Forwarding Engine might get stuck after 100G/400G interface flaps. [PR1638410](#)
- Interoperability issue between legacy line cards and MPC10E/11E might cause incorrect load balancing over aggregate ethernet links. [PR1638489](#)
- [tunnel] [grtag] : mx304 :: [show services inline ip-reassembly statistics fpc 0 pfe-slot 0, not displaying any default output]. [PR1638520](#)
- There is a mismatch between user-configured wavelength and actually transmitted wavelength on 400G-ZR wavelength setting with 75GHz spacing. [PR1638603](#)

- CCL:NGPR: RPD_KRT_RESPONSE_ERROR: krt change failed for prefix <> error from kernel is "EINVAL
-- Bad parameter in request. [PR1638745](#)
- The dynamic tunnel might flap every 15 mins with a non-forwarding route. [PR1639134](#)
- LC9600 having less required planes for fabric bandwidth degradation behaviour. [PR1639212](#)
- JUNOS: JDI_FT_REGRESSION:SUBSCRIBER_SERVICES:MX480: Time difference is not as expected when DUT exports interface-queue-stats to ipfix-collector tool after changing reporting-interval. [PR1639378](#)
- On MX304 device, LC: Temperature value stuck at "Testing" under show chassis fpc output, after powering off PFE-0 & 1. [PR1639602](#)
- On MX304 device, LC: AFT command show sandbox tokens shows output in endless loop, after powering off PFE-0 & 1. [PR1639679](#)
- The show system errors error-id CLI might show inconsistent threshold data. [PR1640264](#)
- pon TLVs from PPPoE-IA tags are not displayed in the show subscribers extensive when preference is set to dsl. [PR1640277](#)
- IPv4 and v6 packet header corruption could happen with some sampling scenario [PR1641119](#)
- KRT queue entries are stuck during Routing Engine switchover when backup RPD is not yet ready. [PR1641297](#)
- Out of order packets might be seen in multicast streams with MVPN extranet scenario, if multicast source route is missing in the receiver VRF. [PR1641323](#)
- The show ldp p2mp tunnel might not display the correct information. [PR1641412](#)
- The show network-agent statistics gnmi detail CLI command is reporting packet drops for some gnmi target-defined mode sensors. [PR1641483](#)
- FPC start time is incorrect under show chassis fpc details CLI command. [PR1641515](#)
- Traffic might be dropped due to the RX queue being full. [PR1641793](#)
- GNMI Capabilities() RPC returns yang module's latest-revision' value instead of "semantic-version"(or oc-ext:openconfig-version) for OpenConfig yang modules. [PR1641936](#)
- [Telemetry] Filtering option for components name(CHASSIS, SIB) fails with /components/component sensor subscription. [PR1641949](#)
- CFM traceoptions writes on every other line. [PR1642948](#)
- On MX480 devices, PFED CPU increased post ISSU and remains around 65 to 75 percentage for 32,000 Layer 2 VPN sBNG services. [PR1643077](#)

- Options to configure VXLAN will not be available under set interfaces FTI unit tunnel encapsulation. [PR1643078](#)
- Traffic Loss might be observed when deactivate or activate the firewall filter. [PR1643187](#)
- Traffic impact might be seen if persistent-learning is enabled on an interface. [PR1643258](#)
- ICMP TTL exceeded packets are not sent out of the switch. [PR1643457](#)
- Post AutoBandwidth Make-before-break, traffic over Conditional metric enabled LSP might get blackholed. [PR1643587](#)
- On MX304 devices, CRC/ALign errors reported on channel 2 and 3 when number of ports is 5 and in pic mode 10G. [PR1643636](#)
- [subscriber_services] [all] : :JDI-REG-SUBSCRIBER_SERVICES:mx480:DHCP:Verify DHCP client count failed and 64,000 DHCPv4 subscribers are not bound as expected count. [PR1643863](#)
- VRRP and IS-IS fails to converge after interface flap. [PR1643932](#)
- RIP NSR state might be stuck in InProgress after Routing-Engine switchover. [PR1644274](#)
- The openconfig network-instance/protocols/static-routes/static id list-key might generate an error on encoding through NETCONF. [PR1644319](#)
- PCEP SRv6 code points changed as per IANA. [PR1644332](#)
- OAMD process is not enabled on MX10008 and MX10016 causing GRE keepalives adjacency down. [PR1644480](#)
- Stateful sync failing between active and backup MX device chassis. [PR1644579](#)
- The video console for vRR might not work after an upgrade to Junos OS with upgraded FreeBSD. [PR1644806](#)
- Traffic drop with EBGP multipath and EBGP paths equal to the maximum-ecmp limit. [PR1645296](#)
- On MX104 devices, **request chassis afeb restart** returns timeout error. [PR1645322](#)
- Issue is seen while bringing up dual stack DHCP subscribers. Not able to bring DHCP subscribers, as subscribers are getting logged out automatically. facing difficulties in RC analysis, as events are received from different daemons. [PR1645574](#)
- The eBGP session might not be established on MX platforms with MS-MPC and MX-SPC3 cards. [PR1645585](#)
- ud and ut pseudo devices stay down after hosting FPC was restarted. [PR1645671](#)
- The alarm might not be generated for EDAC errors until the FPC is rebooted. [PR1646339](#)

- JDI-RCT : Error messages "[Feb 1 10:09:27.665 LOG: Err] stats_lu_counter_read_internal(1430): pfe 0:[NH Cntr] jnh[0x1] != cnh" seen after loading configs [PR1646401](#)
- The show ancp subscriber details CLI output for Access Loop Encapsulation tlv is updated. [PR1647180](#)
- DHCP subscriber traffic might get dropped due to the rpf-check filter. [PR1647214](#)
- Services might not work with the VLAN-rewrite configuration. [PR1647294](#)
- JDI_MX_REGRESSIONS:[MULTICAST]upstream rpf session status in not expected state stuck in Init state.[PR1647746](#)
- Alfaromeo BOOT CPLD version showing 0.0.0. [PR1647823](#)
- When PTP with PHY-timestamping is enabled, significant clock frequency drift might be seen. [PR1647901](#)
- The chassis-control subsystem went to unresponsive state after FHP phase2 trigger. [PR1648030](#)
- Modifying NH that uses indirect nh addr to setting decapsulate_header does not work. [PR1648162](#)
- Subscriber load-balancing not supported in release. [PR1649062](#)
- The afd might crash on MPC10 line cards. [PR1649499](#)
- [subscriber_services/5g-pfcp] [show] mx480 :: Test-11_3 (RLI-49857) fails in CATS. ~5000 sessions in wait state during login after apfe failover. [PR1649861](#)
- The IPv4 traffic drop might be observed in EVPN scenario. [PR1650854](#)
- Chassis alarm **Major CB 0 external-1/0 LOS** can be seen after upgrading to 21.2R3. [PR1651490](#)
- [Illinois:K8] cannot bind subscribers on a UP after the access interface has been disabled and enabled. [PR1652203](#)
- The interface on copper SFP takes 2 times hold-time up timer to come online. [PR1652647](#)
- The rpd might crash when BMP rib-out monitoring is configured for flow-spec route. [PR1653130](#)
- BGP PIC Edge might cause traffic Black-holing after selector corruption. [PR1653562](#)
- Due to the MAC learning limit being exceeded traffic drop might be observed in the MC-AE scenario. [PR1653926](#)
- When fib-streaming is enabled and two or more collectors are involved, fibtd core might be observed due to a timing sync issue. [PR1653942](#)
- The ARP might not resolve with the native-vlan configuration. [PR1654215](#)
- LACP sent IN SYNC to server facing interface when core-isolation is in effect. [PR1654459](#)

- In a subscriber scenario, AAA module of **mobiled** process might cause memory leak on the standby routing-engine. [PR1654947](#)
- BFD might flap when the hold down/up timer is configured. [PR1655088](#)
- [gRIBI]IPIP tunnel remains in the Packet Forwarding Engine even after clearing all programmed route entries. [PR1655531](#)
- Certificate-based VPN tunnel is not established. [PR1655571](#)
- On MX304 device chassis, RSI command execution getting aborted due to error message - "[: JUNIPER: unexpected operator". [PR1655827](#)
- Decap and look up in default instance might not work if backup Nexthop groups are added after adding the routes. [PR1656561](#)
- Delete AFT Operation for PolicyForwrdingEntry is not deleting policy filter for single SINGLE_PRIMARY mode. [PR1657208](#)
- The low priority stream may get stuck and all traffic might be dropped [PR1657378](#)
- Fabric Destination error/Fabric plane in check state. [PR1658164](#)
- The CPU usage SPMB can hit 100% for a short while. [PR1658206](#)
- The rpd crash might be triggered when the BGP route resolves over another BGP route. [PR1658678](#)
- The multipath route might be missing when multipath is configured. [PR1659255](#)
- DHCP relay no-snoop might not work with DHCP local server in the same routing-instance. [PR1613738](#)
- DHCP ALQ needs a new configuration parameter to adjust failover times. [PR1631770](#)
- The rpd-agent crash might be observed once routing processes exit. [PR1637391](#)
- The jdhcpd daemon might crash after Junos OS upgrade. [PR1649638](#)

Infrastructure

- Management port on RE-S-1800X4 might constantly flap after upgrade to 21.2 or higher. [PR1605173](#)
- The lattice-isp tool is much slower and always fail with LTS19 ULC. [PR1607871](#)
- Recovery snapshot might fail if OAM volume is already mounted. [PR1639991](#)

Interfaces and Chassis

- CFM sessions are not up after evo-pfemand restart or gets crashed. [PR1634721](#)
- VRRP route tracking for routes in VRF might not work if **chained-composite-next-hop ingress l3vpn** is used. [PR1635351](#)
- Some daemons might get stuck when snmpd is at 100% CPU utilization. [PR1636093](#)
- [VALE] [USB-Upgrade] JDI_REG_TPTX_REGRESSIONS: The FPCs are not online with USB upgrade from 21.3R1.9 to 21.4R1.11. [PR1637636](#)
- The show vrrp extensive doesn't show the next IFL **Interface VRRP PDU statistics**. [PR1637735](#)
- On Junos OS 20.3 and later release, the tracking routes of VRRP might become unknown after upgradation. [PR1639242](#)
- Traffic loss might be seen for the MAC addresses learned on the ICL interface. [PR1639713](#)
- The aggregate Ethernet interface with 400GE gets flapped on adding or removing a 400GE member link. [PR1641585](#)
- Traffic might be impacted due to the RCP session number reaching the maximum limit. [PR1643855](#)
- The vrrpd core might be observed after interface state change. [PR1646480](#)
- The lacpd might not come up on one of the links in the aggregate Ethernet bundle. [PR1647145](#)
- Authenitcation key can not be configured more than 15 character. [PR1650873](#)
- VRRP failover over 2 seconds might be observed. [PR1652549](#)
- Configuring a VIP from a different subnet (other than parent IP) might affect IPv6 VRRP sessions. [PR1658326](#)
- The MAC address might be learned over the wrong interface in the MC-AE scenario. [PR1658742](#)
- VRRP operations might flap when configuration changes are committed under unrelated VRRP groups present on the same physical interface. [PR1658966](#)

J-Web

- Significant performance improvements were made to JWeb in this release. [PR1652676](#)

Junos XML API and Scripting

- On Junos OS, certificate validation is skipped when fetching system scripts from a HTTPS URL (CVE-2022-22156). [PR1542229](#)

Layer 2 Features

- The Pseudo Wires might go down in VPLS scenario. [PR1655858](#)

Layer 2 Ethernet Services

- Option 82 might not be attached on DHCP request packets. [PR1625604](#)
- IPv6 IA_NA or IA_PD routes might get deleted from the DHCPv6 client. [PR1629171](#)
- Traffic loss might happen when there is a mismatch of subscribers between the primary and backup relay. [PR1638050](#)
- Aggregated Ethernet interface remains up instead of down after deleting loopback and ae interface IP on neighbor while verifying BFD sessions on router. [PR1640240](#)
- The jdhcpd core might be seen if TCP connection is restarted between the ALQ peers. [PR1644919](#)
- DHCP packets might not be sent to the clients when forward-only is reconfigured under the routing instance. [PR1651768](#)

MPLS

- Unified ISIS BFD sessions might take a long time to recover when the interface flaps. [PR1593959](#)
- Standby secondary LSP might be stuck on the same path as primary LSP upon reoptimization. [PR1615326](#)
- LDP protections paths might not be established when auto-targeted-session configuration statement is deactivated and activated. [PR1620262](#)
- Corner case with set protocols mpls label-switched-path <lsp name> no-decrement-ttl when used along with **chained-composite**. [PR1621943](#)
- The timestamp of **MPLS traceroute detail** info doesn't calculate sub-second data properly. [PR1632449](#)
- VCCV BFD session keeps flapping between MX and peer device if ultimate-hop popping is enabled. [PR1634632](#)

- [mpls] [LDP-Tunneling] : mx2020 :: rpd core@ldp_destroy_lib is observed in mx2020 after post Gress. [PR1635863](#)
- The rpd memory leak might be observed in a subscriber management environment with RSVP. [PR1637645](#)
- LSP over broadcast segment stays down when RSVP setup protection is enabled. [PR1638145](#)
- Dynamic bypass LSP might flap at every re-optimization interval. [PR1639292](#)
- When the primary path goes down MPLS LSP does not use the most preferred path after the primary path restoration. [PR1640918](#)
- RPD cores when P2MP Egress interface deleted while LDP p2mp MBB is in progress. [PR1644952](#)
- FRR is not triggered and Backup LSP is not triggered upon IGP change when RSVP graceful-restart/no-reliable is configured without global graceful-restart. [PR1648833](#)
- LSPs which are using the TED Database on Junos OS platforms running BGP LS might not be able to compute paths properly. [PR1650724](#)
- Route stays Up but LSP state is blown off from the downstream node after there is a second failure on the primary LSP link. [PR1654226](#)
- Routing Engine kernel crash might be observed in the one-hop-LSP MPLS scenario with Routing Engine outbound traffic if routing-option resolution preserve-nexthop-hierarchy is configured globally. [PR1654798](#)
- The "rpd" process might get crash when container Label Switch Path (LSP) is configured with default-template. [PR1655177](#)

Network Management and Monitoring

- Configuring set system no-hidden-commands blocks/denies netconf/junoscript sessions. [PR1590350](#)
- Rtsdbd core might be seen when IPsec configuration is activated and deactivated [PR1610594](#)
- The babeltrace core might be triggered in a rare condition. [PR1637992](#)
- VTEP might report a high speed on the sub-interface, causing SNMP alarms. [PR1651774](#)

Platform and Infrastructure

- Error message **gencfg_cfg_msg_gen_handler drop** might be seen after running commit command. [PR1629647](#)
- The packet drop might be seen on FPC on Trio based platforms. [PR1631313](#)

- IP monitor might install default route with incorrect preference value when multiple IP monitoring is configured. [PR1634129](#)
- During ISSU on MX platforms fpc crash might be observed. [PR1637618](#)
- AUTO-CORE-PR : JDI-RCT vRCT : vmxt_lnx core found @ topo_get_link jnh_features_get_jnh jnh_stream_attach. [PR1638166](#)
- The input-vlan-map (pop) might not work on PS interfaces if the native VLAN is in use on the uplink interface. [PR1640254](#)
- Routing Engine switchover might result in traffic loss in a certain scenario. [PR1643416](#)
- The FPC crash might be observed during unified ISSU. [PR1648473](#)
- SCB reset with Error : zfchip_scan line = 844 name = failed due to PIO error.s [PR1648850](#)
- LMEM Parity Error in shared LMEM are not handled properly. [PR1652416](#)
- Regressions : ifstraced.core-tarball.0.tgz found @ inet_ntop6>_inet_ntop>rtplib_ifsm_info_print>dump_all_ifstate_tracerec_from_kmem>ifstraced_go_trace. [PR1654737](#)
- Packet Forwarding Engine might get disabled if a packet with a small size is transmitted out of the queue. [PR1657203](#)
- Lockout-period might not work as expected. [PR1660931](#)

Routing Policy and Firewall Filters

- Existing routing policies might change when global default route-filter walkup is changed. [PR1646603](#)

Routing Protocols

- SHA-1 system login password format not accepted post upgrade. [PR1571179](#)
- The rpd dump file might be seen while processing the BGP updates. [PR1626717](#)
- The traffic might be ceased in PIM scenario. [PR1627990](#)
- The rpd might crash when BGP labeled-unicast family routes are present and BGP multipath is turned on. [PR1630987](#)
- The BGP family route-target might not work in hierarchical Route Reflector scenario. [PR1635018](#)

- The BFD session might be down when multiple addresses of same subnet are configured. [PR1635700](#)
- The multicast traffic might get dropped in the Packet Forwarding Engine. [PR1638141](#)
- On all Junos platforms, when a route is received with broadcast address as next-hop, the platform might not consider it as an invalid route and drop the traffic intended for the destination. [PR1643178](#)
- The BGP peer might stay down in shards after doing a rollback. [PR1643246](#)
- The BGP route might still be present in the multi-path route after increased IGP cost. [PR1643665](#)
- Passive BGP session in no-forwarding instance could not come up. [PR1645010](#)
- Traffic impact might be seen due to failure of IS-IS shortcuts. [PR1645414](#)
- The show multicast snooping route extensive instance evpn-vxlan-A with VLAN filter is not showing VE, AR mesh group route entries. [PR1649410](#)
- Ipv6 Inline BFD sessions are down when neighbor is not resolved. [PR1650677](#)
- Delay in BGP session establishment due to longer time for the listening task to be ready on all platforms running rpd. [PR1651211](#)
- BGP PIC Protection is not working in virtual router. [PR1653356](#)
- RPD core might occur while accessing IFL of mpls-lsp-interfaces in IS-IS (FA-LSP). [PR1654162](#)
- An RPD process crash might be observed, when the received prefix count exceeds configured prefix-limit. [PR1655228](#)
- A policy with a policy action community configuration might not work. [PR1660424](#)

Services Applications

- L2TP tunnels might go down and not able to re-establish after restarting the bbe-smgd process. [PR1629104](#)
- The kmd crash might be observed in IPsec scenario. [PR1637906](#)
- DTCP radius-flow-tap fails to program Packet Forwarding Engine when trigger X-NAS-Port-Id exceeds 48 character length. [PR1647179](#)

Subscriber Access Management

- Event-timestamp in radius Acct-Stop might show future time. [PR1643316](#)
- DHCP clients with static IP addresses binding might get disconnected. [PR1650243](#)

- Pool drain with APM not working. [PR1652715](#)
- JDI-RCT:BBE:Authd core@thr_kill () at thr_kill.S:3. [PR1655832](#)

User Interface and Configuration

- The addition or deletion of the gRPC configuration might cause a memory leak in the EDO application. [PR1619974](#)
- Commit might fail for backup Routing Engine. [PR1636385](#)
- Unable to access configure exclusive mode after mgd process is killed. [PR1641025](#)
- Ignore the syslog - UI_MOTD_PROPAGATE_ERROR: Unable to propagate login announcement (motd) to /var/etc/motd.junos. [PR1642743](#)
- [passive_monitoring] [monitoring] : Scapa : EVO:JDI_FT_REGRESSION: LAZURITE :: Observing configd object-info anomalies at net:juniper::config::interface::IFDCEtherOptionsCommon. [PR1643192](#)

VPNs

- Type 7 routes might be lost in MVPN+PIM SSM scenario. [PR1640487](#)
- The Multicast Tunnel interface is not selected as per the configuration for the Draft-Rosen. [PR1642182](#)
- The routing protocol process might stop working when de-activating and activating the same provider tunnel from one to another instance in a single commit. [PR1647149](#)
- MVPN Inter-AS option B shows updated PMSI attribute tunnel id when advertising type-3 routes to Intra-AS PE. [PR1652481](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading to Release 22.2R1 | [126](#)
- Procedure to Upgrade to Junos OS | [126](#)
- Upgrade and Downgrade Support Policy for Junos OS Releases | [129](#)

- Upgrading a Router with Redundant Routing Engines | 129
- Downgrading from Release 22.2R1 | 130

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.

Basic Procedure for Upgrading to Release 22.2R1



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 Junos OS

To download and install 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-x86-32-22.2R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-x86-64-22.2R1.9-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-x86-32-22.2R1.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-x86-64-22.2R1.9-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**

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 22.2R1, 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: 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](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

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

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

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 8: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Upgrading a Router with Redundant Routing Engines

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

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine, and save the configuration change to both Routing Engines.

2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.
3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Downgrading from Release 22.2R1

To downgrade from Release 22.2R1 to another supported release, follow the procedure for upgrading, but replace the 22.2R1 jinstall package with one that corresponds to the appropriate release.



NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [What's New | 131](#)
- [What's Changed | 132](#)
- [Known Limitations | 132](#)
- [Open Issues | 133](#)
- [Resolved Issues | 134](#)
- [Migration, Upgrade, and Downgrade Instructions | 135](#)

These release notes accompany Junos OS Release 22.2R1 for the NFX Series Network Services Platforms. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [Class of Service | 131](#)
- [Unified Threat Management \(UTM\) | 131](#)
- [Virtualized Network Functions \(VNFs\) | 132](#)

Learn about new features introduced in Junos OS Release 22.2R1 for NFX Series.

Class of Service

- **Support for explicit congestion notification (ECN) (NFX Series, SRX380, SRX300, SRX320, SRX340, SRX345, and vSRX 3.0)**—Starting with Junos OS Release 22.2R1, you can enable ECN marking for packets in scheduler queues on the listed devices. ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. You must enable ECN on both endpoints and on all intermediate devices between the endpoints.

ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets.

To enable ECN, issue the `set class-of-service scheduler-name explicit-congestion-notification` command.

Unified Threat Management (UTM)

- **Web filtering support to nonstandard ports (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 22.2R1, we've extended the Web filtering support for HTTP or HTTPS traffic to nonstandard ports.

[See [web-filtering](#) and [show security utm web-filtering status](#).]

Virtualized Network Functions (VNFs)

- **Support for virtual route reflector (vRR) VNF (NFX250 NextGen)**—In Junos OS Release 22.2R1 and later releases, you can implement the vRR capability on NFX250 NextGen devices by deploying a vRR VNF. Note that single-root I/O virtualization (SR-IOV) interfaces do not support the vRR VNF feature.

This feature is also available in Junos OS Release 21.4R2.

[See [Configuring VNFs on NFX250 NextGen Devices](#).]

What's Changed

There are no changes in behavior and syntax in Junos OS Release 22.2R1 for NFX Series devices.

Known Limitations

IN THIS SECTION

- [Interfaces | 132](#)

Learn about known limitations in this Junos OS Release 22.2R1 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

- When you run the `show interface` command on NFX150 devices to check the interface details, the system does not check whether the interface name provided is valid or invalid. The system does not generate an error message if the interface name is invalid. [PR1306191](#)

Open Issues

IN THIS SECTION

- General Routing | [133](#)
- Interfaces | [133](#)
- Virtual Network Functions (VNFs) | [133](#)

Learn about open issues in Junos OS Release 22.2R1 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.

General Routing

- On the NFX350, if you change the device operational mode to custom mode, ovs-vswitchd cores might be seen on the device. [PR1634245](#)
- On the NFX150 devices, after loading 22.2R1.1, the fablinks go down and the cluster status displays an FL. [PR1664636](#)

Interfaces

- When you run a show interface command on NFX150 devices to check the interface details, the system does not check whether the interface name provided is valid or invalid. The system does not generate an error message if the interface name is invalid. [PR1306191](#)

Virtual Network Functions (VNFs)

- On NFX150 devices, before reusing a VF to Layer 3 data plane interfaces (for example, ge-1/0/3), which was earlier allocated to a VNF, you must restart the system. [PR1512331](#)

- The NFX350 device stops responding after you delete a VNF with SRIOV interfaces. Also, JDM becomes unreachable. As a workaround, you can power cycle the device. [PR1664814](#)

Resolved Issues

IN THIS SECTION

- General Routing | [134](#)
- Virtual Network Functions (VNFs) | [134](#)

Learn about the issues fixed in this release for NFX Series.

General Routing

- On NFX250 device, core files are dumped into the device when you delete vmhost VLANs. [PR1637649](#)

Virtual Network Functions (VNFs)

- On all the NFX devices that have a VNF interface configured with trust mode enabled, VRRP is not functional.

To resolve this issue, you must disable the spoof-check, using the CLI `set virtual-network-functions vnf-name interfaces interface-name mapping interface virtual-function disable-spoof-check`. [PR1643164](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 135](#)
- [Basic Procedure for Upgrading to Release 22.2 | 136](#)

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.

Starting in Junos OS release 21.2R1, all Junos OS products which were previously running on FreeBSD 11.x based Junos OS are migrated to FreeBSD 12.x based Junos OS, except EX4400. Starting with Junos OS release 21.3R1, EX4400 platforms are migrated to FreeBSD 12.x based Junos OS.



NOTE: For information about NFX product compatibility, see [NFX Product Compatibility](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

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

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

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 9: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Basic Procedure for Upgrading to Release 22.2

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



NOTE: The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the device, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the device. For more information, see the [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 22.2R1:

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

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

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

Junos OS Release Notes for PTX Series

IN THIS SECTION

- [What's New | 138](#)
- [What's Changed | 142](#)
- [Known Limitations | 148](#)
- [Open Issues | 149](#)
- [Resolved Issues | 151](#)
- [Migration, Upgrade, and Downgrade Instructions | 153](#)

These release notes accompany Junos OS Release 22.2R1 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

- [Interfaces | 138](#)
- [IP Tunneling | 138](#)
- [Junos Telemetry Interface | 138](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 139](#)
- [Routing Protocols | 140](#)
- [Routing Policy and Firewall Filters | 141](#)
- [Additional Features | 142](#)

Learn about new features introduced in this release for the PTX Series.

Interfaces

- **Custom multipliers for longer LACP hold times (PTX1000, PTX5000, PTX10002, PTX10008, and PTX10016)**—Starting in Junos OS Release 22.2R1, you can configure an LACP hold time that is up to 30 times the periodic interval. When the hold time expires, Junos OS deletes the adjacencies, which results in traffic loss. Use a longer hold time for a peer upgrade to run seamlessly without any traffic loss.

To configure the Junos OS multiplier to calculate the hold time, use the `multiplier` statement at the `[edit interfaces ae aggregated-ether-options lacp]` hierarchy level.

See [\[lacp \(Aggregated Ethernet\)\]](#).

IP Tunneling

- **Sharding support for dynamic IP-over-IP tunneling (MX240, MX480, MX960, PTX1000, PTX10001, and QFX10002)**—Starting in Junos OS Release 22.2R1, we are supporting sharding for dynamic tunnels that are created as a result of BGP route resolution over a tunnel route. BGP uses this tunnel route as a helper route for route resolution.

Junos Telemetry Interface

- **Support for breakout port state sensor (MX240, MX480, MX960, MX2010, MX2020, MX10008, MX10016, and PTX5000)**—Junos OS Release 22.2R1 introduces the `breakout port state / components/component/port/breakout-mode/groups/group/state` for Junos telemetry interface

(JTI) based on the OpenConfig data model **openconfig-platform-port.yang** version 0.4.0. This sensor provides the operational state data for the breakout group identified by the index on platforms running on JUNOS OS.

[See [Telemetry Sensor Explorer](#) and [sensor \(Junos Telemetry Interface\)](#).]

- **Support for CoS sensor (PTX5000)**—Starting in Junos OS Release 22.2R1, use the resource path `/junos/system/linecard/page-drops/page-drop/` to stream CoS page-drop counters and interface details from a device to a collector.

[See [Telemetry Sensor Explorer](#).]

- **Support for CPU state sensor (ACX710, ACX5448, MX204, MX240, MX150, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, and PTX10002)**—Starting in Junos OS Release 22.2R1, use the resource path `/system/cpus/cpu/state/` to export CPU parameters and including CPU usage per process and CPU usage per Routing Engine core information from a device to a collector.

[See [Telemetry Sensor Explorer](#).]

- **Support for forwarding table sensor (MX2020 and PTX5000)**—Junos OS Release 22.2R1 extends support for forwarding information base (FIB) streaming on JTI to include non-default virtual routing and forwarding (VRF) instances.

[See [Telemetry Sensor Explorer](#).]

- **Support for Ethernet interface sensors (MX240, MX480, MX960, MX2010, MX2020, MX10003, MX10008, MX10016, and PTX5000)**—Starting in Junos OS Release 22.2R1, use the subscription path `/interfaces/interface/` or `/interfaces/interface/ethernet/state` to stream Ethernet packet statistics from a device to a collector.

[See [Telemetry Sensor Explorer](#).]

- **Network instance support enhancements (ACX710, ACX5448, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, and PTX10002)**—Starting in Junos OS Release 22.2R1, JTI supports new sensors for network instance statistics for the OpenConfig modules **openconfig-network-instance.yang** and **openconfig-routing-policy.yang**. The support includes OpenConfig configuration and streaming of state data.

[See [Telemetry Sensor Explorer](#) for telemetry support and [OpenConfig User Guide](#) for configuration.]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **BGP classful transport (CT) support for IPv6 and Segment Routing Traffic-Engineered (SR-TE) color-only support (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 22.2R1, we support BGP-CT with IPv6 and BGP service-routes with a color-only mapping community. We have

also enhanced the transport-class configuration statement to provide strict resolution without falling back on best-effort tunnels.

[See [use-transport-class](#), [BGP Classful Transport \(BGP-CT\) with Underlying Colored SR-TE Tunnels Overview](#).]

Routing Protocols

- **DCSPF support for SR-TE with Flex Algo (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, vMX, PTX1000, PTX3000, and PTX5000)**—Starting in Junos OS Release 22.2R1, we support the flexible algorithm (Flex Algo) as a constraint in the compute profile of a segment routing-traffic engineering (SR-TE) LSP. The computation combines any constraints in the compute profile with the ones in the Flex Algo definition to find the resultant path. It uses the Flex Algo segment identifiers (SDIs) in the configuration to compress the resultant path.

We support the feature only for IPv4 SR-MPLS SDIs. You can use SR-TE policy constraints to further fine-tune Flex Algo constraints.

[See [Enabling Distributed CSPF for Segment Routing LSPs](#).]

- **TCP-AO for RPKI validation sessions (MX204, MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2008, MX2010, MX2020, PTX1000, PTX10002, PTX10008, PTX10016, and vRR)**—Starting in Junos OS Release 22.2R1, you can use TCP Authentication Option (TCP-AO) to authenticate resource public key infrastructure (RPKI) validation sessions for securing the Internet's routing infrastructure, such as BGP. Using RPKI, legitimate holders of Internet number resources can control the operation of Internet routing protocols to prevent route hijacking and other attacks.

To enable a TCP-AO chain to authenticate an RPKI validation session, use `authentication-algorithm ao` and the configured authentication-key-chain `keychain` at the `[edit routing-options validation group group_name session address]` and `[edit routing-options validation group group_name]` hierarchy levels.

See [\[TCP Authentication Option \(TCP-AO\)\]](#).

- **Nonstop active routing (NSR) support with BGP RIP sharding and BGP UpdateIO features (ACX5048, ACX5096, ACX5448, MX240, MX960, MX2008, MX10016, and PTX5000)**—Starting in Junos OS Release 22.2R1, we've enabled nonstop routing (NSR) for BGP RIP sharding and BGP UpdateIO features. With NSR enabled, the backup Routing Engine and backup routing protocol process (rpd) become the primary Routing Engine without negatively affecting the BGP peering sessions with the neighbors if the primary Routing Engine fails. The backup rpd processes the replicated BGP control-plane information and populates the route state in the same multithreaded manner as in the primary rpd.

After you configure NSR, the `show bgp neighbor` and `show bgp summary` commands display the information about the specific shards in the backup Routing Engine. To display the replicated information for a specific shard in the `show bgp replication` command, use the `rib-sharding shard-name` option.

See [[show bgp neighbor](#), [show bgp summary](#), [show bgp replication](#), and [BGP Overview](#).]

- **BGP extended route retention (MX960, PTX1000, and QFX10002)**—In Junos OS Release 22.2R1, we've enhanced the long-lived graceful restart (LLGR) capabilities for a BGP helper device. With this feature enabled, Junos OS supports LLGR helper mode regardless of the BGP peer LLGR capabilities. We've introduced a new configuration statement `extended-route-retention` at the `[edit protocols bgp group neighbor graceful-restart long-lived]` hierarchy level. We've also updated the outputs of the following operational commands:

- `show bgp neighbor`
- `show route extensive`

[See [graceful-restart-long-lived-edit-protocols-bgp](#).]

- **Anomaly checker for rpd object reference count (MX Series, PTX Series, and QFX Series)**—In Junos OS Release 22.2R1, we introduce a generic reference count infrastructure that all the modules in `rpd` can use. The module maintains lock and unlock statistics corresponding to each object type in use. Any application can call the `refcount` increment or decrement API when an object is referred. The module also provides a mechanism to detect anomalies such as a leak or overflow in an object's `refcount`.
- **Origin validation communities conversion to keywords (MX10008 and PTX10016)**—Starting in Junos OS Release 22.2R1, you can choose to accept or reject the origin validation extended communities received from an eBGP peer. The default behavior of Origin Validation State Extended Community (OVS EC) changes to *rejected* if the extended community is received from an eBGP peer. You can configure your device to accept the community when needed. We also support the configuration of distinguished communities with keywords (valid, invalid, and unknown) at all the three layers of the BGP configuration hierarchy—global, group, and per-neighbor. If you enable the OVS EC at a hierarchy level, it's enabled for the lower levels as well. However, you can choose to disable it explicitly at a lower layer if required at any instance.

Routing Policy and Firewall Filters

- **Support for multiple named validation databases from multiple sources (MX204 and PTX10016)**—Starting in Junos OS Release 22.2R1, we support multiple named validation databases from multiple sources. You can also consult validation databases across instances and track RIBs that consult the various databases to enable notification when entries are modified.

To Specify a named route-validation database, use `validation-state (invalid | valid)` option at the `[edit routing-options validation database <database-name> static record <destination> maximum-length <prefix-length> origin-autonomous-system <as-number>]` hierarchy level.

To Specify target route-validation database for a validation session, use `database <database-name>` option at the `[edit routing-options validation group <group-name> session]` hierarchy level.

To specify validation database, use validation-database-instance option at the [edit policy-statement <policy-name> term <term-name> from] hierarchy level.

[See [policy-statement](#), [session \(Origin Validation for BGP\)](#), and [validation \(Origin Validation for BGP\)](#).]

Additional Features

Support for the following features has been extended to these platforms.

- **Collect ON_CHANGE BGP RIB telemetry statistics and BGP neighbor telemetry with sharding (MX Series, PTX Series and QFX Series)**

[See [Telemetry Sensor Explorer](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 142](#)
- [General Routing | 143](#)
- [Junos XML API and Scripting | 144](#)
- [MPLS | 144](#)
- [Network Management and Monitoring | 144](#)
- [Routing Protocols | 146](#)
- [User Interface and Configuration | 146](#)
- [VPNs | 148](#)

Learn about what changed in this release for PTX Series.

Authentication and Access Control

- **SHA-1 password format deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX and vSRX)**—We've removed the sha1 option at the [edit system login password format] hierarchy level because SHA-1 is no longer supported for plain-text password encryption.

General Routing

- Change in in unnumbered-address support for GRE tunnel?Starting in Junos OS Release 24.4R1, there is a behavioural change in unnumbered-address support for GRE tunnel with IPV6 family and display donor interface for both IPV4 and IPV6 families of GRE tunnel. You can view interface donor details under show interfaces hierarchy level. **See show interfaces.** <https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/command/show-interfaces-gigabit-ethernet.html>
- OpenConfig container names for Point-to-Multipoint per interface ingress and egress sensors are modified for consistency from "signalling" to "signaling".
- **Enhancement to snmp mib command behavior (PTX10008)**—Starting in Junos OS Evolved, when you execute show snmp mib walk decimal command, the output parameter jnxRedundancySwitchoverReason is not working as expected, which always show the value 0 instead of expected values. Now, jnxRedundancySwitchoverReason output parameter is corrected to expected behavior with the following expected values.

jnxRedundancySwitchoverReason OBJECT-TYPE SYNTAX INTEGER { other(1), -- others neverSwitched(2), -- never switched userSwitched(3), -- user-initiated switchover autoSwitched(4) -- automatic switchover }

[See [show snmp mib](#).]

- **No support for PKI operational mode commands on the Junos Limited version (MX Series routers, PTX Series routers, and SRX Series devices)**—We do not support request, show, and clear PKI-related operational commands on the limited encryption Junos image ("Junos Limited"). If you try to execute PKI operational commands on a limited encryption Junos image, then an appropriate error message is displayed. The pkid process does not run on Junos Limited version image. Hence, the limited version does not support any PKI-related operation.
- **The <request-system-zeroize/> RPC response indicates when the device successfully initiates the requested operation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the <request-system-zeroize/> RPC successfully initiates the zeroize operation, the device emits the <system-zeroize-status>zeroizing re0</system-zeroize-status> response tag to indicate that the process has started. If the device fails to initiate the zeroize operation, the device does not emit the <system-zeroize-status> response tag.
- **Support for Embedded RP on PTX10008**—From this release, we support the Embedded RP feature on PTX10008 devices.

[See [Configuring Embedded RP](#).]

- **"Switchover Status Ready" incorrectly describes the status of the backup Routing Engine (RE) after node reboot (PTX10004, PTX10008, PTX10016)**—During preparation for switchover between master RE and backup RE running Junos OS Evolved releases prior to 22.2R1, "Switchover Status Ready" from the show system switchover command on the backup RE node, after system reboot,

incorrectly describes the status of the backup RE. The incorrect status description results from a discrepancy between the master RE and the backup RE both using local uptime to determine if sufficient time had elapsed before declaring "Switchover Status Ready".

Use the `request chassis routing-engine master switch` command on the master RE and the backup RE to obtain the correct status when preparing for switchover.

[See [show system switchover](#) and [request chassis routing-engine master](#).]

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is operational mode command, include the `cert-file` option and `authentic`. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation. When you refresh a script using the `request system scripts refresh-from` specify the certificate path. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

See [request system scripts refresh-from](#).

See [cert-file](#).

MPLS

- When defining a constrained path LSP using more than one strict hop belonging to the egress node, the first strict hop must be set to match the IP address assigned to the egress node on the interface that receives the RSVP Path message. If the incoming RSVP Path message arrives on an interface with a different IP address the LSP is rejected.

Network Management and Monitoring

- **Limits increased for the `max-datasize` statement (ACX Series, PTX Series, and QFX Series)**—The `max-datasize` statement's minimum configurable value is increased from 23,068,672 bytes (22 MB) to 268,435,456 bytes (256 MB), and the maximum configurable value is increased from 1,073,741,824

(1 GB) to 2,147,483,648 (2 GB) for all script types. Furthermore, if you do not configure the `max-datasize` statement for a given script type, the default maximum memory allocated to the data segment portion of a script is increased to 1024 MB. Higher limits ensure that the device allocates a sufficient amount of memory to run the affected scripts.

[See [max-datasize](#).]

- **Changes when deactivating or deleting instances of the ephemeral configuration database (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The following changes apply when you deactivate or delete ephemeral database instances in the static configuration database:
 - When you deactivate the entire `[edit system configuration-database ephemeral]` hierarchy level, the device deletes the files and corresponding configuration data for all user-defined ephemeral instances. In earlier releases, the files and configuration data are preserved; however, the configuration data is not merged with the static configuration database.
 - When you delete an ephemeral instance in the static configuration database, the instance's configuration files are also deleted. In earlier releases, the configuration files are preserved.
 - You can delete the files and corresponding configuration data for the default ephemeral database instance by configuring the `delete-ephemeral-default` statement in conjunction with the `ignore-ephemeral-default` statement at the `[edit system configuration-database ephemeral]` hierarchy level.

[See [Enable and Configure Instances of the Ephemeral Configuration Database](#).]

- **Support for automatically synchronizing an ephemeral instance configuration upon committing the instance (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, QFX Series, and vMX)**—You can configure an ephemeral database instance to synchronize its configuration to the other Routing Engine every time you commit the ephemeral instance on a dual Routing Engine device or an MX Series Virtual Chassis. To automatically synchronize the instance when you commit it, include the `synchronize` statement at the `[edit system commit]` hierarchy level in the ephemeral instance's configuration.

[See [Commit and Synchronize Ephemeral Configuration Data Using the NETCONF or Junos XML Protocol](#).]

- **Changes to the NETCONF `[edit-config]` RPC response (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the `[edit-config]` operation returns an error, the NETCONF server does not emit a `load-error-count` element in the RPC response. In earlier releases, the `[edit-config]` RPC response includes the `load-error-count` element when the operation fails.
- **DES deprecation for SNMPv3**—The Data Encryption Standard (DES) privacy protocol for SNMPv3 is deprecated due to weak security and vulnerability to cryptographic attacks. For enhanced security, configure the triple Data Encryption Standard (3DES) or the Advanced Encryption Standard (CFB128-AES-128 Privacy Protocol) as the encryption algorithm for SNMPv3 users.

[See [privacy-3des](#) and [privacy-aes128](#).]

Routing Protocols

- **The RPD OSPF LDP SYNC message not logged**—On all Junos OS and Junos OS Evolved devices, when an LDP session goes down there is a loss of synchronization between LDP and OSPF. After the loss of synchronization, when an interface has been in the holddown state for more than three minutes, the system log message with a warning level is sent. This message appears in both the messages file and the trace file. However, the system log message does not get logged if you explicitly configure the hold-time for ldp-synchronization at the [edit protocols ospf area area id interface interface name] hierarchy level less than three minutes. The message is printed after three minutes.
- To achieve consistency among resource paths, the resource path /mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counterip-addr='address'/state/countersname='name'/out-pkts/ is changed to /mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counterip-addr='address'/state/countersname='name/'. The leaf "out-pkts" is removed from the end of the path, and "signalling" is changed to "signaling" (with one "I").
- When the krt-nexthop-ack statement is configured, the RPD will wait for the next hop to get acknowledged by PFE before using it for a route. Currently, only BGP-labeled routes and RSVP routes support this statement. All other routes will ignore this statement.
- **SSH TCP forwarding disabled by default**—We've disabled the SSH TCP forwarding feature by default to enhance security. To enable the SSH TCP forwarding feature, you can configure the allow-tcp-forwarding statement at the edit system services ssh hierarchy level.

In addition, we've deprecated the tcp-forwarding and no-tcp-forwarding statements at the edit system services ssh hierarchy level.

[See [services \(System Services\)](#).]

User Interface and Configuration

- **Load JSON configuration data with unordered list entries (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos schema requires that list keys precede any other siblings within a list entry and appear in the order specified by the schema. Junos devices provide two options to load JSON configuration data that contains unordered list entries:

- Use the `request system convert-json-configuration` operational mode command to produce JSON configuration data with ordered list entries before loading the data on the device.
- Configure the `reorder-list-keys` statement at the `[edit system configuration input format json]` hierarchy level. After you configure the statement, you can load JSON configuration data with unordered list entries, and the device reorders the list keys as required by the Junos schema during the load operation.
- When you configure the `reorder-list-keys` statement, the load operation can take significantly longer to parse the configuration, depending on the size of the configuration and number of lists. Therefore, for large configurations or configurations with many lists, we recommend using the `request system convert-json-configuration` command instead of the `reorder-list-keys` statement.

[See [json](#) and [request system convert-json-configuration](#)]

- **Junos XML protocol Perl modules deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—We no longer provide the Junos XML protocol Perl client for download. To use Perl to manage Junos devices, use the NETCONF Perl library instead.

[See [Understanding the NETCONF Perl Client and Sample Scripts..](#)]

- A new field `rollback pending` is added to the output of `show system commit` that identifies whether `commit confirmed` is issued. It is removed once `commit` or `commit check` is issued or `commit confirmed` is rolled back after `rollback timeout`.
- When you configure `max-cli-sessions` at the `[edit system]` hierarchy level, it restricts the maximum number of CLI sessions that can coexist at any time. Once the `max-cli-sessions` number is reached, new CLI access is denied. The users who are configured to get the CLI upon login, are also denied new login. The `max-cli-sessions` is configured so you can control the memory usage for the CLI. You may set the `max-cli-sessions` per your requirement. However, if `max-cli-sessions` is not configured, there is no control on the number of CLIs getting invoked.
- Persistent CLI timestamps—To have a persistent CLI timestamp for the user currently logged in, enable the `set cli timestamp` operational command. This ensures the timestamp shows persistently for each new line of each SSH session for the user or class until the configuration is removed. To enable timestamp for a particular class with permissions and format for different users, configure the following statements: `set system login class class name permissions permissions`, `set system login class class name CLI timestamp`, `set system login user username class class name authentication plain-text-password`



NOTE: The default timestamp format is `%b %d %T`. You can modify the format per your requirements. For example, you can configure the following statement: `set system login class class name CLI timestamp format "%T %b %d`. To enable timestamp for a particular user with default class permissions and format, configure the following statements: `set system`

```
login user username class class name authentication plain-text-password and set system login
user username CLI timestamp.
```

VPNs

- **Changes to show mvpn c-multicast and show mvpn instance outputs**— The FwdNh output field displays the multicast tunnel (mt) interface in the case of Protocol Independent Multicast (PIM) tunnels.

[See [show mvpn c-multicast](#).]

Known Limitations

IN THIS SECTION

- [MPLS | 148](#)

Learn about known limitations in this release for the PTX Series.

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

MPLS

- The automatic sorting of configuration entries does not work if defined under group. [PR1637730](#)

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 149](#)
- [General Routing | 149](#)
- [MPLS | 150](#)
- [Multicast | 151](#)

Learn about open issues in Junos OS Release 22.2R1 for the PTX Series.

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

Class of Service (CoS)

- The class-of-service (CoS) values associated with default code-point aliases are inconsistent with Junos OS. Therefore, the incoming packets associates with incorrect CoS servicing level (forwarding class and packet loss priority (PLP)). [PR1667404](#)

General Routing

- On the PTX Series routers with FPC-PTX-P1-A or FPC2-PTX-P1A, you might encounter a single event upset (SEU) event that might cause a linked-list corruption of the TQCHIP. The following syslog message gets reported: "Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt_min_free_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero Jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002." The Junos OS Chassis Management error handling detects such a condition, raises an alarm, and disables the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, restart the FPC. Contact your Juniper support representative if the issue persists even after the FPC restarts. [PR1254415](#)

- On routers and switches running Junos OS, with Link Aggregation Control Protocol (LACP) enabled, deactivating a remote aggregated Ethernet (AE) member link makes the local member link move to LACP detached state and cause traffic drops on that member link. The same scenario applied when a new member link is added where the other end of that link is not yet configured with LACP.

[PR1423707](#)

- In link aggregation groups (LAGs) enhanced link aggregation group (a.k.a. aggregated Ethernet child NH reduction/AE simple/LAG enhanced) scenario, if FPC hosting single child member aggregated Ethernet bundle is rebooted, the aggregate next-hops on such aggregated Ethernet bundles might be discarding traffic, as the logical interfaces reroute tables might get stuck in the down state, the packets drop might be seen on all other FPCs that the ingress traffic towards such aggregated Ethernet bundles and the aggregated Ethernet interfaces might not be used. [PR1551736](#)

- Copying files to /tmp/ causes a huge JTASK_SCHED_SLIP. Copy files to /var/tmp/ instead. [PR1571214](#)

- In PTX Series router under MAC statistics output-mac-control-frames and output-mac-pause-frames does not increment. [PR1610745](#)

- Tunnel statistics displays incorrect values because of the tunnel interfaces cache flow. [PR1627713](#)

- When sending BGP-LU traffic or Layer 3 VPN traffic over IPIP tunnels, if the remote end device is an IP that does not understand labels. The labeled unicast or Layer 3 VPN label cannot go on top.

[PR1631671](#)

- V6 default route will not get added after successful DHCPv6 client binding on PTX1000 router during ZTP. [PR1649576](#)

- Junos OS Release 20.2R1 and later has introduced firmware version for PHY chip on some MIC and PIC models. After that migration, some of the 100GE ports on MIC/PIC or its peer devices might see PCS and framing errors. [PR1651526](#)

- On all Junos OS PTX Series platforms, an FPC heap memory leak occurs while creating the composite multicast next-hop. [PR1661286](#)

MPLS

- With the rsvp local reversion configuration a PLR originates the "Bw_unavailable PathErr" during Fast Reroute (FRR). Junos Label Edge Router (LER or ingress router) ignores this type of PathErr message. However, this can be a problem if an ingress LER implementation reacts to this PathErr by bringing down the protected LSP causing packet loss. [PR1670638](#)

Multicast

- On Junos OS PTX Series platforms, the traffic might silently drop because of the next-hop installation failure for multicast Resource Reservation Protocol (RSVP) Point to Multipoint (P2MP) traffic. This failure might encounter in a scaled RSVP P2MP environment after a network event that cause reconvergence.[PR1653920](#)

Resolved Issues

IN THIS SECTION

- General Routing | [151](#)
- Interfaces and Chassis | [152](#)
- Layer 2 Ethernet Services | [152](#)
- MPLS | [153](#)
- Routing Protocols | [153](#)
- User Interface and Configuration | [153](#)

Learn about the issues fixed in this release for PTX Series.

General Routing

- As part of ON-CHANGE initial synchronization data we do not export the logical interfaces state counters. [PR1620160](#)
- QSFP in slot et-0/0/0 might not come up after plugging-in. [PR1620527](#)
- The mcontrol might frequently miss keepalives from backup Routing Engine. [PR1624623](#)
- SNMP trap message for FPC restart shows FRU removal instead of FRU offline or FRU power off. [PR1629738](#)
- Multiple link flaps and traffic loss might occur. [PR1630006](#)
- ON_CHANGE telemetry does not work for "backplane-facing-capacity" sensors. [PR1635606](#)

- SPMB might crash immediately after a switchover. [PR1637950](#)
- KRT change fails for prefix *variable* error from kernel will display "EINVAL -- Bad parameter in request". [PR1638745](#)
- KRT queue entries get stuck during Routing Engine switchover when backup RPD is not ready. [PR1641297](#)
- Filtering option for component name (CHASSIS, SIB) fails with /components/component sensor subscription. [PR1641949](#)
- Junos OS: RIB and Packet Forwarding Engines can get out of synchronization because of the memory leak during interface flap or route churn (CVE-2022-22209). [PR1642172](#)
- Enabling traffic over the conditional metric, LSP might silently drop or discard. [PR1643587](#)
- After rebooting FPC, the EDAC errors might not generate alarm. [PR1646339](#)
- mac-vrf does not support MAC limit configuration. [PR1647327](#)
- BGP sensor "/bgp-rib/afi-safis/afi-safi/ipv4-unicast/loc-rib/" not available as a 'periodic' sensor. [PR1649529](#)
- FPCs might restart unexpectedly upon receipt of specific MPLS packets with certain multi-unit interface configurations (CVE-2022-22202). [PR1649586](#)
- IRP memory parity issue might result in traffic loss on the PTX Series routers. [PR1650217](#)
- The traffic with ether type 0X88FC might corrupt. [PR1651703](#)
- Configuring gre-key in firewall filter might break the DSCP classification. [PR1652762](#)
- On PTX Series platforms in segment routing IPv6 (SRv6) scenario, traffic might drop when you configure END.DT46 and END.DT4. [PR1655518](#)

Interfaces and Chassis

- The FPCs won't come up on upgrading USB from Junos OS Release 21.3R1.9 to Junos OS Release 21.4R1.11. [PR1637636](#)

Layer 2 Ethernet Services

- Traffic loss might be seen on QFX10000 due to congestion. [PR1635935](#)

- Aggregated Ethernet interface remains up after deleting loopback and aggregated Ethernet interface IP on neighbor while verifying BFD sessions on router. [PR1640240](#)

MPLS

- IS-IS BFD sessions may take a long time to recover when the interface flap. [PR1593959](#)

Routing Protocols

- The rpd process might generate a core file while processing the BGP updates. [PR1626717](#)
- Denial of Service (DoS) vulnerability in RPD upon receipt of specific BGP update (CVE-2022-22213). [PR1642741](#)
- The BGP route might be present in the multi-path route after increasing IGP cost. [PR1643665](#)

User Interface and Configuration

- Unable to access configure exclusive mode after killing mgd process. [PR1641025](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading to Release 22.2 | [154](#)
- Upgrade and Downgrade Support Policy for Junos OS Releases | [156](#)
- Upgrading a Router with Redundant Routing Engines | [157](#)

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.

Starting in Junos OS release 21.2R1, all Junos OS products which were previously running on FreeBSD 11.x based Junos OS are migrated to FreeBSD 12.x based Junos OS, except EX4400. Starting with Junos OS release 21.3R1, EX4400 platforms are migrated to FreeBSD 12.x based Junos OS.

Basic Procedure for Upgrading to Release 22.2

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



NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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



NOTE: The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the router, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [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 22.2R1:

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-x86-64-22.2R1.9.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-x86-64-22.2R1.9-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 22.2 `jinstall` package, you cannot return to the previously installed software by issuing the `request system software rollback` command. Instead, you must issue the `request system software add validate` command and specify the `jinstall` package that corresponds to the previously installed software.

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

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

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

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

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases.

Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 10: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Upgrading a Router with Redundant Routing Engines

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

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

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for QFX Series

IN THIS SECTION

- [What's New | 158](#)
- [What's Changed | 164](#)
- [Known Limitations | 170](#)
- [Open Issues | 171](#)
- [Resolved Issues | 175](#)
- [Migration, Upgrade, and Downgrade Instructions | 181](#)

These release notes accompany Junos OS Release 22.2R1 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

- [Class of Service | 159](#)
- [EVPN | 159](#)
- [IP Tunneling | 161](#)
- [MACsec | 161](#)
- [Network Management and Monitoring | 162](#)
- [Routing Protocols | 162](#)
- [Routing Policy and Firewall Filters | 162](#)
- [Additional Features | 163](#)

Learn about new features introduced in this release for QFX Series switches.

Class of Service

- **Support for additional IEEE 802.1 TLVs (QFX5110, QFX5120-48Y, and EX4650)**—Starting in Junos OS Release 22.2R1, QFX5110, QFX5120-48Y, and EX4650 switches support the IEEE 802.1 TLVs listed below. The TLVs advertise information to peers, and the Link Layer Discovery Protocol (LLDP) neighbors use the TLVs to discover a device's capabilities.
 - Link aggregation TLV—IEEE 802.1 subtype 7
 - Enhanced transmission selection (ETS) configuration TLV—IEEE 802.1 subtype 9
 - ETS recommendation TLV—IEEE 802.1 subtype A
 - Priority-based flow control (PFC) configuration TLV—IEEE 802.1 subtype B

[See [show lldp neighbors](#) and [show dcbx neighbors](#).]

EVPN

- **Support for blocking asymmetric EVPN Type 5 routes (MX960, QFX5110, and QFX10002)**—Starting in Junos OS Release 22.2R1, you can configure the local node to reject asymmetric EVPN Type 5 routes on EVPN-VXLAN networks. The local node examines the incoming EVPN Type 5 route packets and rejects the route when the virtual network identifier (VNI) in the ingress route differs from the locally configured VNI.

To block asymmetric EVPN Type 5 routes, include the `reject-asymmetric-vni` statement at the `[edit routing-instance routing-instance-name protocols evpn ip-prefix-routes]` hierarchy level.

[See [EVPN Type 5 Route with VXLAN encapsulation for EVPN-VXLAN](#) and [ip-prefix-routes](#).]

- **Automatically derived ESI configuration (MX Series, QFX5100, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX10002, QFX10002-60C, QFX10008, and QFX10016)**—In the current implementation, Junos OS derives the Ethernet segment identifier (ESI) from the system ID and the administrative key on the local multihomed provider edge (PE) device that is a part of the LACP link (actor). Starting in Junos OS Release 22.2R1, you can also configure the multihomed devices on an EVPN-VXLAN network to automatically generate the ESI from:

- The system ID and administrative key on the remote customer edge (CE) device (partner).
- The locally configured `mac` and local discriminator values.

To automatically derive the ESI using the system ID and administrative key on the remote CE device, include `type-1-lacp` at the `[edit interfaces aeX aggregated-ether-options lacp auto-derive]` hierarchy level.

To automatically derive the ESI using locally configured values, configure `mac` and `local-discriminator` at the `[edit interfaces aeX aggregated-ether-options lacp auto-derive type-3-system-mac]` hierarchy level.

[See [Understanding Automatically Generated ESIs in EVPN Networks](#).]

- **EVPN/VXLAN MAC filtering and transit VNI match support for pure IPv6 underlay (QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.2R1, we support MAC filtering on a Layer 2 interface in the EVPN-VXLAN context. We've also implemented VXLAN network identifier (VNI) matching on source and destination IP outer headers for transit traffic on a Layer 3 interface. The device matches VNI values on outer headers only, and on ingress traffic only. On transit devices that are routing tunnel packets, MAC filtering must support matching the VNI in the outer header, along with outer header source and destination IPv6 addresses as match conditions. Use the VNI match filter under the `vxlan match` CLI option for the `set firewall family inet6 filter` term from `vxlan vni vni-id` command. Use the `show firewall filter` command to display statistics.

[See [MAC Filtering, Storm Control, and Port Mirroring Support in an EVPN VXLAN Environment](#).]

- **Support for service provider style CLI in EVPN-VXLAN Layer 3 gateways (EX4650, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 22.2R1, you can use the service provider style CLI to configure a Layer 3 gateway in an edge-routed bridging scenario. In this scenario, you can map an IRB interface to a virtual network identifier and perform VXLAN routing.

You can use the service provider CLI when the interface VLAN ID is the same or different from the VLANs VLAN-ID. If the VLAN ID is different, the VLAN ID can be “none” or between 1 and 4,000.

[See [EVPN User Guide](#).]

- **Optimized intersubnet multicast (OISM) with MAC-VRF instances and IGMPv2 or IGMPv3 in an EVPN-VXLAN fabric (EX4650, QFX5110, QFX5120, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.2R1, you can configure OISM on leaf devices and border leaf devices in an EVPN-VXLAN ERB overlay fabric with:
 - MAC-VRF routing instances or the default switch instance with IGMPv2 or IGMPv3.
 - IGMP snooping and selective multicast Ethernet tag (SMET) forwarding optimizations with IGMPv2 or IGMPv3.

When you configure OISM, you must enable OISM and IGMP snooping on all the server leaf and border leaf devices in the EVPN-VXLAN fabric. With a MAC-VRF instance configuration, you configure the OISM supplemental bridge domain (SBD) and all revenue VLANs in the MAC-VRF instances on all leaf and border leaf devices in the fabric.

[See [Optimized Intersubnet Multicast in EVPN Networks](#).]

- **Assisted replication (AR) integrated with optimized intersubnet multicast (OISM) in an EVPN-VXLAN ERB fabric (QFX5110, QFX5120, QFX10002-32Q, QFX10002-72Q, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.2R1, you can configure AR and OISM together in an EVPN-VXLAN ERB overlay fabric.

You can configure the following devices in each AR role in an integrated AR and OISM environment:

- AR replicator: QFX10002-32Q, QFX10002-72Q, QFX10008, and QFX10016
- AR leaf: QFX5110, QFX5120, QFX10002-32Q, QFX10002-72Q, QFX10008, and QFX10016

Here is a summary of integrated AR and OISM support:

- AR leaf devices can be OISM server leaf or border leaf devices.
- AR replicator devices can operate in either collocated mode (the device is both an AR replicator and an OISM border leaf device) or standalone mode (the device is an AR replicator but not an OISM border leaf or server leaf device). In ERB fabrics, a standalone mode AR replicator is usually a lean spine device.
- AR replicator devices must be running Junos OS software that supports OISM (even when operating in standalone mode).

When you configure AR devices:

- You can configure the EVPN instances using the default switch instance or using MAC-VRF instances (with `vlan-based` or `vlan-aware` service types only).
- With standalone mode, you must configure the AR replicator devices with the same tenant VRF instances, corresponding IRB interfaces, and member VLANs as the OISM border leaf and server leaf devices.

[See [Assisted Replication Multicast Optimization in EVPN Networks](#) and [Optimized Intersubnet Multicast in EVPN Networks](#).]

IP Tunneling

- **Sharding support for dynamic IP-over-IP tunneling (MX240, MX480, MX960, PTX1000, PTX10001, and QFX10002)**—Starting in Junos OS Release 22.2R1, we are supporting sharding for dynamic tunnels that are created as a result of BGP route resolution over a tunnel route. BGP uses this tunnel route as a helper route for route resolution.

MACsec

- **Certificate-based authentication and encryption for MACsec (MX Series)**—Starting in Junos OS Release 22.2R1, you can enable MACsec on links connecting switches or routers using certificate-based authentication and encryption. Connected devices can mutually authenticate using 802.1X over Extensible Authentication Protocol-Transport Layer Security (EAP-TLS) and dynamically derive the connectivity association key (CAK) for encryption.

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

Network Management and Monitoring

- **sFlow support for EVPN-VXLAN (QFX10002-60C, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 22.1R1, the QFX10000 line of switches set up for EVPN-VXLAN with an IPv4 underlay and overlay support sFlow monitoring technology.

You can view the **Extended router data** and **Extended switch data** headers on the collectors as part of the sFlow records.

The **Extended switch data** header contains information of the Format, Flow data length (byte), Next hop, Next hop source mask, and Next hop destination mask fields.

The **Extended router data** header contains information of the Format, Flow data length (byte), Incoming 802.1Q VLAN, Incoming 802.1p priority, Outgoing 802.1Q VLAN, and Outgoing 802.1p priority fields.

[See [Overview of sFlow Technology](#).]

Routing Protocols

- **BGP extended route retention (MX960, PTX1000, and QFX10002)**—In Junos OS Release 22.2R1, we've enhanced the long-lived graceful restart (LLGR) capabilities for a BGP helper device. With this feature enabled, Junos OS supports LLGR helper mode regardless of the BGP peer LLGR capabilities. We've introduced a new configuration statement `extended-route-retention` at the `[edit protocols bgp group neighbor graceful-restart long-lived]` hierarchy level. We've also updated the outputs of the following operational commands:

- `show bgp neighbor`
- `show route extensive`

[See [graceful-restart-long-lived-edit-protocols-bgp](#).]

- **Anomaly checker for rpd object reference count (MX Series, PTX Series, and QFX Series)**—In Junos OS Release 22.2R1, we introduce a generic reference count infrastructure that all the modules in rpd can use. The module maintains lock and unlock statistics corresponding to each object type in use. Any application can call the `refcount` increment or decrement API when an object is referred. The module also provides a mechanism to detect anomalies such as a leak or overflow in an object's `refcount`.

Routing Policy and Firewall Filters

- **Support for firewall filters per logical interface (QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200, and QFX5210)**—Starting in Junos OS Release 22.2R1, you can configure port firewall filters per logical interface, in the input direction, using the service provider-style configuration. To configure, use the `set chassis per-logical-interface-firewall` CLI command. In earlier Junos OS releases, port firewall filters would be applied to all logical interfaces of a physical interface.

- **Optimize TCAM when EVPN/VXLAN is enabled (EX4400-48F, EX4650, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200, and QFX5210)—**

In Junos OS Release 22.2R1, we've introduced CLI configuration commands to optimize ternary content addressable memory (TCAM) space usage. Use these commands to prevent ingress filter processor (IFP) TCAM space exhaustion:

- `set chassis ivacl-firewall-no-portrange-profile`
- `set chassis iracl-firewall-ipv4-profile`
- `set chassis ipvacl-firewall-l2-profile`
- `set chassis input-firewall-optimized-profile`

Additional Features

Support for the following features has been extended to these platforms.

- **Collect ON_CHANGE BGP RIB telemetry statistics and BGP neighbor telemetry with sharding (MX Series, PTX Series and QFX Series)**

[See [Telemetry Sensor Explorer](#).]

- **Inband Flow Analyzer (IFA) 2.0 (QFX5120-48YM and QFX5120-48T)—**We've extended support for IFA 2.0 to the QFX5120-48YM and QFX5120-48T switches. You can now configure the MTU and maximum clip length for IFA packets on all QFX5120 switches. You can also set the IFA clock source on the QFX5120-48YM switch.

[See [Inband Flow Analyzer \(IFA\) 2.0 Probe for Real-Time Performance Monitoring](#).]

- **Support for BPDU protection for EVPN-VXLAN (EX4300-48MP, EX4400, EX4650, QFX5110, QFX5120, QFX10002-60C, QFX10008, and QFX10016)—**Starting in Junos OS Release 22.2R1, we support bridge protocol data unit (BPDU) protection for EVPN-VXLAN.

[See [Understanding BPDU Protection for EVPN-VXLAN](#).]

- **Support for flow-based telemetry (QFX5120).** You can configure flow-based telemetry (FBT) for VXLAN-encapsulated traffic only. FBT for VXLANs enables per-flow-level analytics for VXLAN-encapsulated traffic that uses either a centrally routed bridging (CRB) overlay or an edge-routed bridging (ERB) overlay. FBT for VXLANs uses inline monitoring services to create flows, collect them, and export them to a collector.

FBT for VXLANs differs from the current FBT feature in the following ways:

- Only IRB interfaces are supported.
- Only one inline-monitoring instance and one collector are supported.

- The flow table keys include fields important for VXLANs, and the IPv4 and IPv6 flow table keys include different fields.
- You cannot configure an option template identifier or a forwarding class.

[See [Flow-Based Telemetry for VXLANs](#).]

- **Supported transceivers, optical interfaces, and DAC cables** Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.
- **Symmetric integrated routing and bridging (IRB) with EVPN Type 2 routes** (EX4400, EX4650, EX9204, EX9208, EX9214, MX Series, vMX, QFX5110, QFX5120, QFX10002, QFX10002-60C, QFX10008, and QFX10016). We support this feature only with MAC-VRF EVPN routing instance configurations and MAC-VRF service types `vlan-based` and `vlan-aware`. [See [Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics](#) and [irb-symmetric-routing](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 165](#)
- [General Routing | 165](#)
- [Junos XML API and Scripting | 165](#)
- [MPLS | 166](#)
- [Network Management and Monitoring | 166](#)
- [Routing Protocols | 168](#)
- [User Interface and Configuration | 168](#)
- [VPNs | 170](#)

Learn about what changed in this release for QFX Series.

Authentication and Access Control

- **SHA-1 password format deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX and vSRX)**—We've removed the sha1 option at the [edit system login password format] hierarchy level because SHA-1 is no longer supported for plain-text password encryption.

General Routing

- **The request-system-zeroize RPC response indicates when the device successfully initiates the requested operation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the request-system-zeroize RPC successfully initiates the zeroize operation, the device emits the <system-zeroize-status>zeroizing re0</system-zeroize-status> response tag to indicate that the process has started. If the device fails to initiate the zeroize operation, the device does not emit the system-zeroize-status response tag.
- OpenConfig container names for Point-to-Multipoint per interface ingress and egress sensors are modified for consistency from "signalling" to "signaling".
- **Instance type change is not permitted from default to L3VRF in open configuration (ACX Series, EX Series, MX Series, QFX Series, SRX Series, vMX, and vSRX)**—DEFAULT_INSTANCE is the primary instance that runs when there is no specific instance type configured in the route set routing-options. Any instance you explicitly configure is translated into set routing-instance r1 routing-options. The issue appears in translation, when you change instance type DEFAULT_INSTANCE (any instance to DEFAULT_INSTANCE) to L3VRF or L3VRF to DEFAULT_INSTANCE. As a result, such changes are not permitted. Additionally, DEFAULT_INSTANCE can only be named DEFAULT, and DEFAULT is reserved for DEFAULT_INSTANCE, therefore allowing no such changes.

Junos XML API and Scripting

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

configure the cert-file statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

See [request system scripts refresh-from](#).

See [cert-file](#).

MPLS

- When defining a constrained path LSP using more than one strict hop belonging to the egress node, the first strict hop must be set to match the IP address assigned to the egress node on the interface that receives the RSVP Path message. If the incoming RSVP Path message arrives on an interface with a different IP address the LSP is rejected.

Network Management and Monitoring

- **Enhanced system log messages (SRX Series, NFX Series, and QFX5130, QFX5200, QFX5220, and QFX5700)**—We've added multiple events inside the event tag using the UI_LOGIN_EVENT|UI_LOGOUT_EVENT format, which has an option (|) to separate the events, to generate system log messages.

Earlier to this release, the event tag used the UI_LOGIN_EVENT UI_LOGOUT_EVENT format and for various combinations of <get-syslog-events> rpc filters was not getting logged.

[See [Overview of System Logging](#).]

- **Limits increased for the max-datasize statement (ACX Series, PTX Series, and QFX Series)**—The max-datasize statement's minimum configurable value is increased from 23,068,672 bytes (22 MB) to 268,435,456 bytes (256 MB), and the maximum configurable value is increased from 1,073,741,824 (1 GB) to 2,147,483,648 (2 GB) for all script types. Furthermore, if you do not configure the max-datasize statement for a given script type, the default maximum memory allocated to the data segment portion of a script is increased to 1024 MB. Higher limits ensure that the device allocates a sufficient amount of memory to run the affected scripts.

[See [max-datasize](#).]

- **DES deprecation for SNMPv3**—The Data Encryption Standard (DES) privacy protocol for SNMPv3 is deprecated due to weak security and vulnerability to cryptographic attacks. For enhanced security, configure the triple Data Encryption Standard (3DES) or the Advanced Encryption Standard (CFB128-AES-128 Privacy Protocol) as the encryption algorithm for SNMPv3 users.

[See [privacy-3des](#) and [privacy-aes128](#).]

- **Changes when deactivating or deleting instances of the ephemeral configuration database (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The following changes apply when you deactivate or delete ephemeral database instances in the static configuration database:
 - When you deactivate the entire [edit system configuration-database ephemeral] hierarchy level, the device deletes the files and corresponding configuration data for all user-defined ephemeral instances. In earlier releases, the files and configuration data are preserved; however, the configuration data is not merged with the static configuration database.
 - When you delete an ephemeral instance in the static configuration database, the instance's configuration files are also deleted. In earlier releases, the configuration files are preserved.
 - You can delete the files and corresponding configuration data for the default ephemeral database instance by configuring the delete-ephemeral-default statement in conjunction with the ignore-ephemeral-default statement at the [edit system configuration-database ephemeral] hierarchy level.

[See [Enable and Configure Instances of the Ephemeral Configuration Database](#).]

- **Limits increased for the max-datasize statement (ACX Series, PTX Series, and QFX Series)**—The max-datasize statement's minimum configurable value is increased from 23,068,672 bytes (22 MB) to 268,435,456 bytes (256Â MB), and the maximum configurable value is increased from 1,073,741,824 (1Â GB) to 2,147,483,648 (2Â GB) for all script types. Furthermore, if you do not configure the max-datasize statement for a given script type, the default maximum memory allocated to the data segment portion of a script is increased to 1024Â MB. Higher limits ensure that the device allocates a sufficient amount of memory to run the affected scripts.

[See [max-datasize](#).]

- **Support for automatically synchronizing an ephemeral instance configuration upon committing the instance (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, QFX Series, and vMX)**—You can configure an ephemeral database instance to synchronize its configuration to the other Routing Engine every time you commit the ephemeral instance on a dual Routing Engine device or an MX Series Virtual Chassis. To automatically synchronize the instance when you commit it, include the synchronize statement at the [edit system commit] hierarchy level in the ephemeral instance's configuration.

[See [Commit and Synchronize Ephemeral Configuration Data Using the NETCONF or Junos XML Protocol](#).]

- **Changes to the NETCONF [edit-config] RPC response (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the [edit-config] operation returns an error, the NETCONF server does not emit a load-error-count element in the RPC response. In earlier releases, the [edit-config] RPC response includes the load-error-count element when the operation fails.

Routing Protocols

- **The RPD OSPF_LDP_SYNC message not logged**—On all Junos OS and Junos OS Evolved devices, when an LDP session goes down there is a loss of synchronization between LDP and OSPF. After the loss of synchronization, when an interface has been in the holddown state for more than three minutes, the system log message with a warning level is sent. This message appears in both the messages file and the trace file. However, the system log message does not get logged if you explicitly configure the hold-time for ldp-synchronization at the [edit protocols ospf area area id interface interface name] hierarchy level less than three minutes. The message is printed after three minutes.
- To achieve consistency among resource paths, the resource path /mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counterip-addr='address'/state/countersname='name'/out-pkts/ is changed to /mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counterip-addr='address'/state/countersname='name/'. The leaf "out-pkts" is removed from the end of the path, and "signalling" is changed to "signaling" (with one "I").
- When the krt-nexthop-ack statement is configured, the RPD will wait for the next hop to get acknowledged by PFE before using it for a route. Currently, only BGP-labeled routes and RSVP routes support this statement. All other routes will ignore this statement.
- **SSH TCP forwarding disabled by default**—We've disabled the SSH TCP forwarding feature by default to enhance security. To enable the SSH TCP forwarding feature, you can configure the allow-tcp-forwarding statement at the [edit system services ssh] hierarchy level. In addition, we've deprecated the tcp-forwarding and no-tcp-forwarding statements at the [edit system services ssh] hierarchy level.

[See [services \(System Services\)](#)]

User Interface and Configuration

- **Load JSON configuration data with unordered list entries (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos schema requires that list keys precede any other siblings within a list entry and appear in the order specified by the schema. Junos devices provide two options to load JSON configuration data that contains unordered list entries:
 - Use the request system convert-json-configuration operational mode command to produce JSON configuration data with ordered list entries before loading the data on the device.
 - Configure the reorder-list-keys statement at the [edit system configuration input format json] hierarchy level. After you configure the statement, you can load JSON configuration data with unordered list entries, and the device reorders the list keys as required by the Junos schema during the load operation.

- When you configure the `reorder-list-keys` statement, the load operation can take significantly longer to parse the configuration, depending on the size of the configuration and number of lists. Therefore, for large configurations or configurations with many lists, we recommend using the `request system convert-json-configuration` command instead of the `reorder-list-keys` statement.

[See [json](#) and [request system convert-json-configuration](#)]

- Junos XML protocol Perl modules deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—We no longer provide the Junos XML protocol Perl client for download. To use Perl to manage Junos devices, use the NETCONF Perl library instead.

[See [Understanding the NETCONF Perl Client and Sample Scripts](#).]

- A new field `rollback pending` is added to the output of `show system commit` that identifies whether `commit confirmed` is issued. It is removed once `commit` or `commit check` is issued or `commit confirmed` is rolled back after `rollback timeout`.
- When you configure `max-cli-sessions` at the `[edit system]` hierarchy level, it restricts the maximum number of CLI sessions that can coexist at any time. Once the `max-cli-sessions` number is reached, new CLI access is denied. The users who are configured to get the CLI upon login, are also denied new login. The `max-cli-sessions` is configured so you can control the memory usage for the CLI. You may set the `max-cli-sessions` per your requirement. However, if `max-cli-sessions` is not configured, there is no control on the number of CLIs getting invoked.
- Persistent CLI timestamps**—To have a persistent CLI timestamp for the user currently logged in, enable the `set cli timestamp` operational command. This ensures the timestamp shows persistently for each new line of each SSH session for the user or class until the configuration is removed. To enable timestamp for a particular class with permissions and format for different users, configure the following statements:

```
set system login class class name permissions permissions, set system login class class name cli timestamp, and
set system login user username class class name authentication plain-text-password
```



NOTE: The default timestamp format is `%b %d %T`. You can modify the format per your requirements. For example, you can configure the following statement:

To enable timestamp for a particular user with default class permissions and format, configure the following statements:
`set system login user username class <variable>class name</variable> authentication plain-text-password set system login user <variable>username</variable> cli timestamp`

VPNs

- **Changes to show mvpn c-multicast and show mvpn instance outputs**—The FwdNh output field displays the multicast tunnel (mt) interface in the case of Protocol Independent Multicast (PIM) tunnels.

[See [show mvpn c-multicast](#).]

Known Limitations

IN THIS SECTION

- [General Routing | 170](#)

Learn about known limitations in Junos OS Release 22.2R1 for QFX Series switches.

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

General Routing

- On QFX5000, in EVPN_VXLAN deployment, BUM traffic replication over VTEP might send out more packets than expected. [PR1570689](#)
- Unified ISSU on QFX5120-48Y and EX4650 switches will not be supported if there is a change in the Cancun versions of the chipset SDKs between the releases. This is a product limitation as change in the Cancun firmware leads to the chip reset and hence ISSU is impacted. The Cancun versions in the chipset SDKs should be the same between two JUNOS OS releases for ISSU to work. [PR1634695](#)
- The incoming VLAN tag is removed at ingress. So, it is not available at Egress Sampling. [PR1654879](#)
- For traffic dropped at egress due to split horizon in BCM during egress path processing, statistics are shown on vtep as statistics are fetched at ingress pipeline in BCM. [PR1656400](#)
- When VNI ranges span across two beta blocks programming ranges have limitations. The forwarding pipeline places the VNI ID along with the flags in the gre key. As this is specific to implementation and discrete nature of non-overlapping ranges, we do not recommend valid VNI

range; As a workaround, configure VNI matches as specific number OR a list of VNIs instead of a range.[PR1660623](#)

Open Issues

IN THIS SECTION

- [EVPN | 171](#)
- [General Routing | 171](#)
- [Infrastructure | 174](#)
- [Layer 2 Ethernet Services | 174](#)
- [Layer 2 Features | 174](#)
- [Platform and Infrastructure | 175](#)
- [Routing Protocols | 175](#)

Learn about open issues in Junos OS Release 22.2R1 for QFX Series switches.

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

EVPN

- On Junos OS platforms such as MX Series, QFX5000, and QFX10000 platforms and Junos OS Evolved platforms such as ACX Series and MX Series, on interface up and down event loop prevention might not work resulting in BUM traffic might drop. [PR1669811](#)

General Routing

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

- VXLAN VNI (multicast learning) scaling on QFX5110 traffic issue is seen from VXLAN tunnel to Layer 2 interface. [PR1462548](#)
- 5M DAC connected between QFX10002-60C and MX2010 doesn't link up. Interoperability issues occur between QFX10002-60C and MX2010. [PR1555955](#)
- To avoid the additional interface flap, configure interface hold time. [PR1562857](#)
- On a QFX5120, when you disable a protected link, you might see a delay in the system response. It might take 200 to 400 milli seconds for the system to react to disable link event. [PR1579931](#)
- In a fully loaded device, at times, firewall programming fails due to scaled prefix configuration with more than 64,800 entries. [PR1581767](#)
- Pim VxLAN is not working on TD3 chipsets enabling VxLAN flexflow after Junos OS Release 21.3R1. [PR1597276](#)
- On QFX5100, you can see the optical power after detaching and attaching QSFP on disabled interface. [PR1606003](#)
- In QFX10002-60C under mac statistics output-mac-control-frames and output-mac-pause-frames does not increment. [PR1610745](#)
- On QFX5120-48Y, the traffic or the protocols on pure Layer 3 interfaces might behave unexpectedly. This occurs when you load scaled and baseline configurations multiple times one after the other without a pause in between. [PR1612973](#)
- On Junos OS QFX10000 platforms with scaled number of Bidirectional Forwarding Detection (BFD) sessions configured, addition of a new BFD session might cause flapping in newly added session and other existing BFD sessions. [PR1621976](#)
- Backup FPC lose their connection to the master when new members are added to the Virtual Chassis fabric (VCF). [PR1634533](#)
- We don't support the bound delay configuration feature for the logical interface. You will see the core file only when you enable bound delay configuration on the device. [PR1634941](#)
- On Junos platforms, high BGP scale with flapping route and BGP monitoring protocol (BMP) collector/station is slow. Because of the memory pressure the rpd process crashes. [PR1635143](#)
- On QFX10002-60c, in EVPN/VxLAN scenario multicast traffic received on the INET interface (L3 interface) might be dropped. [PR1636842](#)
- When a 100G interface on a QFX5120 converts to a Virtual Chassis port, the interface stays down as the port is configured as 40G internally. [PR1638156](#)

- On the QFX5100 Virtual Chassis platform, although you configure local-minimum-links statement on the device, the aggregated Ethernet interface do not go down on a single interface flap in the bundle. [PR1649637](#)
- You might observe a traffic loss with virtual-router function on QFX5000 line of switches. [PR1650335](#)
- On QFX Series platforms, for a filter change operation, transit packets might drop on the Packet Forwarding Engine. ICMP, BFD and all routing protocols might go down when the timing issue occurs. [PR1651546](#)
- BFD session session-state is showing DOWN when operating in centralized mode. Packet loss might be seen. [PR1658317](#)
- After converting access side port from SP style to EP style, MAC-IP learning fails for a host and ARP do not get resolved. [PR1658657](#)
- On all Junos OS platforms, when a secondary Precision Time Protocol (PTP) member that used one IPv6 address reboots with another IP address via Dynamic Host Configuration Protocol (DHCP). That might end up in multiple entries for the same source port. [PR1659453](#)
- On QFX10000 Junos OS platforms, the configuration of the IGMP group range might result in a specific multicast route get programmed and this might cause traffic loss. [PR1659732](#)
- EX4600 and QFX5100-24Q devices Virtual-chassis will be in unstable state for 3-7 minutes causing traffic loss. [PR1661349](#)
- After upgrading on QFX5110, the IPv6 ND packets drop. As a result, you might not see functional impact on IPv4. [PR1662707](#) [PR1662707](#)
- On QFX5000 Series platforms, when unicast Address Resolution Protocol (ARP) is received for a MAC address that is already learned in an EVPN-VxLAN environment, the ARP request is flooded and duplicate packets might be seen on leaf devices. You might see some service impact where split-horizon might not work or continuous mac-move might be seen. This issue rarely in a production environment due to presence of intermediate switches which might resolve the unicast ARP query. [PR1665306](#)
- After rebooting, you'll see that the link does not have a static MAC address and results in unicast traffic flooding. You might observe a traffic drop for this static MAC. [PR1666399](#)
- NSSU fails with below combinations on QFX5120-48T.
 1. From Junos OS Release 22.1R1.10 to 22.2R1.6
 2. From Junos OS Release 21.4R1.12 to 22.2R1.6
 3. From Junos OS Release 21.3R2.11 to 22.2R1.6.

PR1669702

- On QFX5110 platforms, the native VLAN feature might not be working as expected, as the interface even if configured with native VLAN, the traffic egresses out with tagged VLAN instead of untagged traffic.[PR1669857](#)
- Native vlan with vlan rewrite on same interface not supported..[PR1671372](#)
- In case of EVPN VXLAN IPv6 Underlay, ECN bits from the IP TOS of CE packet is not getting copied to the Tunnel header(Access to Network case). [PR1672308](#)
- On QFX5K Junos devices, bulk configs add/delete for CE facing ports may cause PFE core.[PR1672537](#)

Infrastructure

- On QFX Series line of switches, IPv6 traffic output byte of `ipv6-transit-statistics` is not in range as per traffic generator statistics. [PR1653671](#)

Layer 2 Ethernet Services

- The DHCP client configuration is coming from AIU script and vsdk sandbox. The DHCP client configuration coming from AIU script has the serial Id in vendor id where as the default configuration from sandbox doesn't have. There is no impact on functionality or service. [PR1601504](#)

Layer 2 Features

- In case of access-side interfaces used as SP-style interfaces, when a new logical interface is added and if there is already a logical interface on the physical interface, there is a traffic drop on the existing logical interface. [PR1367488](#)
- On QFX5100 platforms, a change in TPID in the Device Control daemon might result in traffic drop in Packet Forwarding Engine because of the failure in Layer 2 learning or interfaces flapping. [PR1477156](#)

Platform and Infrastructure

- If you use the source address NTP configuration parameter and issue the command `set ntp date` from the CLI, packets will be sent with the source address of the outgoing interface rather than the manually configured IP address. Typically the manually configured IP address would be a loopback address. The problem does not apply to automatically generated NTP poll packets. [PR1545022](#)

Routing Protocols

- On all Junos OS platforms and all Junos OS Evolved platforms, Routing Process Daemon (rpd) crashes and restarts when Border Gateway Protocol (BGP) is configured and a specific timing condition is hit for secondary route. This issue might cause a traffic impact. [PR1659441](#)

Resolved Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 176](#)
- [EVPN | 176](#)
- [General Routing | 176](#)
- [Interfaces and Chassis | 180](#)
- [Layer 2 Ethernet Services | 180](#)
- [MPLS | 180](#)
- [Network Management and Monitoring | 180](#)
- [Platform and Infrastructure | 180](#)
- [Routing Protocols | 181](#)
- [User Interface and Configuration | 181](#)

Learn about the issues fixed in this release for QFX Series.

Class of Service (CoS)

- The uplink interface remains down for a longer duration due to VXLAN scaled configuration [PR1631448](#)

EVPN

- Few ARP/ND/MAC entries for Vlans are missing with MAC-VRF configuration [PR1609322](#)
- The MAC address might not be visible in the EVPN/VXLAN environment [PR1645591](#)
- On all Junos and Junos OS Evolved platforms configured with EVPN-VXLAN, if an Ethernet Segment Identifier (ESI) link towards a multihomed CE experiences a quick down->up->down transition on a leaf, a remote Ethernet VPN (EVPN) peer might continue to install the Virtual Extensible LAN (VXLAN) tunnel (vtep) towards the affected leaf or the failed device in its ESI next-hop, resulting in traffic drops on the leaf. [PR1648368](#)
- On all Junos and EVO platforms, Integrated routing and bridging (IRB) static entry might be missing in the mac-ip-table for MAC move allowed for local static IRB entry. [PR1650202](#)

General Routing

- The dcpfe might crash on QFX5k devices [PR1588704](#)
- During FRR, when more than one multi home interface is down, traffic may loop for QFX5110(Merus) [PR1596589](#)
- Error message "error: syntax error: request-package-validate" will be seen on device cli output during Non Stop Software Upgrade [PR1596955](#)
- QFX5200: Observed dcpfe core-dump while testing ISSU from 21.1R1.11 to 21.2R1.7 [PR1600807](#)
- Chassisd generates "Cannot read hw.chassis.startup_time value: m" every 5 seconds on qfx10008 and qfx10016 [PR1603588](#)
- The packet drop might be seen when packet size exceeds 9000 MTU [PR1615447](#)
- The BFD session might flap on the QFX5120-48YM platform [PR1616692](#)
- One-time interface flap might be seen on the QFX5120 platform [PR1618891](#)

- BGP session may not establish between loopback interfaces when routes are learnt through type5 EVPN routes [PR1620642](#)
- Host generated IPv4 traffic sent over IPv6 next-hop with IRB interface might get dropped [PR1623262](#)
- Interface on QFX52xx not coming up after swapping from 100G to 40G [PR1623283](#)
- PKID could crash and generate a core-file when there was limited memory available on the routing-engine [PR1624613](#)
- The third 802.1Q tag might not be pushed onto the stack in the Q-in-Q tunneling [PR1626011](#)
- Traffic loss might be observed due to Address Resolution Protocol (ARP) getting programmed as indirect next-hop by control plane for external routes distributed over IRB on QFX10k Junos platforms [PR1627876](#)
- DHCP inform ack might be sent with broadcast address when DHCP smart relay is used [PR1628837](#)
- Some ports (port 20 and above) on QFX5110-32Q VC may not come up after a device restart or PFE reboot [PR1629231](#)
- The interface on the peer device might remain up even after disabling the 10G interface on the Juniper device [PR1629637](#)
- Traffic might get dropped when "family ethernet-switching" is configured on the interface in Q-in-Q scenario [PR1629680](#)
- LACP timeout might be observed during high CPU utilization [PR1630201](#)
- QFX5130 as Border Leaf with L3 interface connectivity, traffic forwarding is not happening to Multihomed receiver connected to Border Leafs when the L3 interface goes down on DF side [PR1631249](#)
- Inner VLAN might be stripped off when input-native-vlan-push is disabled [PR1631771](#)
- The interface might remain in the "UP/UP" state even the interface is admin disabled [PR1632440](#)
- You may see a slow response or timeout on the CLI or SNMP with accessing to sxe-0/0/0 on QFX5120-48T-6c. [PR1632620](#)
- The FBF filtered VLAN traffic will not be passed properly to the forwarding routing instances over AE interfaces on QFX5K/EX4600/EX4650 platforms [PR1633452](#)
- The VCPs connected with the AOC cable might not come up after upgrading to 17.3 or later releases [PR1633998](#)

- An EVPN-VXLAN fabric consisting of QFX5000 devices experience high convergence time for few seconds during link down or node reboot events. [PR1634415](#)
- On EX4650-48Y and QFX5120-48Y switches when there is a link flap, the link or interface might take more than expected time to come-up; resulting in increased traffic loss. [PR1634495](#)
- If you configure chassis disk-partition, chassisd might crash. [PR1635812](#)
- Routes might be slow to install in the LPM table. [PR1635887](#)
- You might observe a traffic drop when STP is configured in VxLAN environment. [PR1636950](#)
- The Packet Forwarding Engine might crash while removing the port from a VLAN. [PR1637013](#)
- Configuring L2PT on a transit switch in a Q-in-Q environment breaks L2PT for other S-VLANs. [PR1637249](#)
- MACsec traffic drop might be seen when a back-to-back graceful switchover is performed. [PR1637822](#)
- The interfaces might delay to come up after reboot or transceiver replacement. [PR1638045](#)
- Targeted broadcast or WOL feature might not work on the QFX5000 platforms. [PR1638619](#)
- MAC-move might be observed when dhcp-security is configured [PR1639926](#)
- MAC address of the hosts might get learned on incorrect VLAN which may lead to traffic loss [PR1639938](#)
- On QFX series base license is missing post upgrade to 20.3 and later releases [PR1640123](#)
- A missing release of memory after effective lifetime vulnerability in the kernel of Junos OS allows an unauthenticated network based attacker to cause a Denial of Service (DoS). Refer to <https://kb.juniper.net/JSA69713> for more information. [PR1642172](#)
- ICMP TTL exceeded packets are not sent out of the switch [PR1643457](#)
- Packets are dropped in ingress QFX5K with EVPN-LAG multihoming due to VP-LAG programming issue [PR1644152](#)
- On Junos OS QFX5000 platforms, the LACP packet might get intercepted and the interface detached from LACP when configuring VLAN tagging in the EVPN-VXLAN scenario. [PR1645929](#)
- Vxlan Tunnel termination due to change in config [PR1646489](#)
- The firewall might drop inbound packets if the filter is configured under IRB interface [PR1646740](#)
- OSPF control packets may get dropped due to the "flow check" function in the interoperability case [PR1648272](#)

- On QFX10000 platforms with EVPN-VXLAN is deployed, differences in the unicast and multicast packets between IS-IS hello packets and Connectionless Network Protocol (CLNP) ping packets might be observed. The packet might get egressed out, but the Logical Link Control (LLC) header data is missed. A header needs to be added for all LLC packets, so the peer node cannot recognize these packets without the header, and it simply discards the CLNP packet.[PR1648078](#)
- In EVPN-VxLAN environment, non-VxLAN traffic might be dropped if VxLAN and non-VxLAN traffic share the same ECMP next-hop. [PR1649841](#)
- L2PT on a transit switches like EX4600 and QFX5100 Series in a Q-in-Q environment breaks L2PT for other S-VLANs if uplink is aggregated Ethernet interface.[PR1650416](#)
- The local-bias might stop working after the device is rebooted [PR1651151](#)
- On all Junos OS platforms, with the L2 circuit configuration when the ether type in the incoming packet is 0x88FC is not added to the known ether type list in ASIC for which traffic might get corrupted. [PR1651703](#)
- On all QFX5120 and EX4650 Junos OS platforms, when priority bits (P-bits) in C-tag are copied to S-tag in the switch, and then the C-tag is reset to 0. [PR1652976](#)
- On all QFX10002-36Q, QFX10002-72Q, QFX10008, and QFX10016 platforms, when Service Provider (SP) style is configured with IRB and native-vlan leads to ARP failure. [PR1654215](#)
- LACP sent IN SYNC to server facing interface when core-isolation is in effect. [PR1654459](#)
- The MAC limiting port security feature might not work. [PR1659873](#)
- On Junos OS QFX5200-32C-32Q platform, on disabling and then enabling interfaces, the Fast Ethernet channel (FEC) might mismatch and the FEC details link might not come up.[PR1657534](#)
- Valid software licenses might not be in sync between members in the Virtual chassis. [PR1658913](#)
- On all Junos OS platforms, the Virtual Extensible LAN (VxLAN) port goes to Spanning Tree Protocol (STP) BLK state when the port is down and takes a while to go back to the FWD state when the link is up. Traffic loss is seen in this case. [PR1659533](#)
- QFX Series line of devices are using IPv4 ID field for a OSPF flow check function, but it is violating RFC6864. [PR1660369](#)
- BUM traffic received on CE interface will loopback to ingress interface after removing EVPN VXLAN FRR configuration (reroute-address). [PR1662515](#)
- L2 multicast traffic loss occurs on QFX5120-48T Virtual chassis platform on Junos OS Release 22.2R1.9. [PR1663102](#)
- Adaptive [load balancing] statistics are not updated under show interfaces aeX extensive on QFX10000 platform.[PR1663881](#)

Interfaces and Chassis

- show vrrp extensive doesnot show the next logical interface "Interface VRRP PDU statistics".
[PR1637735](#)
- Traffic loss might be seen for the mac addresses learned on the ICL interface. [PR1639713](#)
- On all Junos platforms, there might be an issue with ARP learning between ICL interface and local MC-AE interfaces when an MC-AE device is misconfigured and the configuration is rolledback using rollback command. This issue might result in packet drop.[PR1648271](#)
- In all Junos platforms, IPv6 neighbor might be learned over the local MC-AE (multi-chassis aggregate ethernet) interface which is in the down state instead of the inter chassis link (ICL).[PR1658742](#)

Layer 2 Ethernet Services

- Traffic loss might be seen on QFX10000 due to congestion. [PR1635935](#)
- Aggregated Ethernet child interfaces with LACP configurations are not timing out even if peer is gone and not sending any bridge protocol data unit (BPDU).[PR1640240](#)

MPLS

- IS-IS BFD sessions might take a long time to recover when the interface flap. [PR1593959](#)

Network Management and Monitoring

- VTEP might report a high speed on the sub-interface, causing SNMP alarms. [PR1651774](#)

Platform and Infrastructure

- The packet drop might be seen on FPC on Trio based platforms. [PR1631313](#)

Routing Protocols

- PIM accept-remote-source knob config removal. [PR1593283](#)
- The BFD session might be down when multiple addresses of same subnet are configured. [PR1635700](#)
- Ipv6 Inline BFD sessions are down when neighbor is not resolved. [PR1650677](#)
- On all Junos and Junos Evolved platforms, when a policy with policy action "community set/add/delete" is configured and the same policy is imported under Routing Information Protocol (RIP) protocol, community configuration might not work. [PR1660424](#)

User Interface and Configuration

- Unable to access configure exclusive mode after mgd process gets killed. [PR1641025](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrading Software on QFX Series Switches | [182](#)
- Installing the Software on QFX10002-60C Switches | [183](#)
- Installing the Software on QFX10002 Switches | [184](#)
- 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 | [185](#)
- Installing the Software on QFX10008 and QFX10016 Switches | [187](#)
- Performing a Unified ISSU | [191](#)
- Preparing the Switch for Software Installation | [191](#)
- Upgrading the Software Using Unified ISSU | [192](#)
- Upgrade and Downgrade Support Policy for Junos OS Releases | [194](#)

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.

Starting in Junos OS release 21.2R1, all Junos OS products which were previously running on FreeBSD 11.x based Junos OS are migrated to FreeBSD 12.x based Junos OS, except EX4400. Starting with Junos OS release 21.3R1, EX4400 platforms are migrated to FreeBSD 12.x based Junos OS.

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 22.2 in the Release pull-down list to the right of the Software tab on the Download Software page.
4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 22.2 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new `jinstall` package on the device.



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

Customers in the United States and Canada use the following command:

```
user@host> request system software add source/jinstall-host-qfx-5-x86-64-22.2R1.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 22.2 jinstall package, you can issue the `request system software rollback` command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a **junos-vmhost-install-x.tgz**.

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot. If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.



NOTE: The QFX10002-60C switch supports only the 64-bit version of Junos OS.



NOTE: If you have important files in directories other than /config and /var, copy the files to a secure location before upgrading. The files under /config and /var (except /var/etc) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

If the installation package resides locally on the switch, execute the **request vmhost software add <pathname><source>** command.

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-22.2R1.9.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-x86-64-22.2R1.9.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-x86-64-22.2R1.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-x86-64-22.2R1.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-x86-64-22.2R1.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-x86-64-22.2R1.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 191](#)
- ["Upgrading the Software Using Unified ISSU" on page 192](#)

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 EX Series Switches](#) 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-x86-64-22.2-R1.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

We have two types of releases, EOL and EEOL:

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

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

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 11: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [What's New | 196](#)
- [What's Changed | 201](#)
- [Known Limitations | 204](#)
- [Open Issues | 204](#)
- [Resolved Issues | 206](#)
- [Migration, Upgrade, and Downgrade Instructions | 211](#)

These release notes accompany Junos OS Release 22.2R1 for the SRX Series Services Gateways. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [Class of Service | 196](#)
- [Flow-Based and Packet-Based Processing | 196](#)
- [High Availability | 197](#)
- [J-Web | 197](#)
- [Juniper Advanced Threat Prevention Cloud \(ATP Cloud\) | 198](#)
- [MPLS | 198](#)
- [Multicast | 199](#)
- [Network Address Translation \(NAT\) | 199](#)
- [Unified Threat Management \(UTM\) | 199](#)
- [VPNs | 199](#)
- [Additional Features | 200](#)

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

Class of Service

- **Support for explicit congestion notification (ECN) (NFX Series, SRX380, SRX300, SRX320, SRX340, SRX345, and vSRX 3.0)**—Starting with Junos OS Release 22.2R1, you can enable ECN marking for packets in scheduler queues on the listed devices. ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. You must enable ECN on both endpoints and on all intermediate devices between the endpoints.

ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets.

To enable ECN, issue the `set class-of-service scheduler-name explicit-congestion-notification` command.

Flow-Based and Packet-Based Processing

- **Support for IPv6 tunnel (SRX Series and vSRX 3.0)**— Starting in Junos OS Release 22.2R1, you can encapsulate IPv4 and IPv6 traffic over the IPv6 network.

The IPv6 tunnel helps IPv4 traffic traverse over the IPv6 network. You can use IPv6 tunneling in various features such as policy routing and preferential billing. For example, a set-top box that supports only IPv4 traffic can traverse the server over an IPv6 network.

[See [show security flow session](#).]

- **NP-cache scale-up (SRX4600)**—Starting in Junos OS Release 22.2R1, the NP-cache wing count is 20 million. With this increment, the number of Express Path sessions increase fourfold.

[See [Sessions per Wing Statistics](#).]

- **Bypass IP block fragmentation check with allowlist configuration (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 22.2R1, you can configure an allowlist for an IP block fragment screen. The traffic from source addresses in the allowlist groups bypasses the IP block fragmentation check. The IP block fragment allowlist supports both IPv4 and IPv6 addresses and a maximum of 32 allowlist groups.

[See [ids.option](#) and [Understanding Allowlists for IP Block Fragment Screen](#).]

High Availability

- **Insert additional SPC3 in a multinode high availability (HA) setup (SRX5800, SRX5600, and SRX5400)**—Starting in Junos OS Release 22.2R1, you can insert SRX5K-SPC3 cards on the SRX5000 line of devices in a multinode HA setup without interrupting traffic.

You can use the new procedure to insert SPC3 cards when the setup has an encrypted inter-chassis link (ICL) or a non-encrypted ICL.

In case of an encrypted ICL setup, use the new statement `set chassis high-availability hardware-upgrade` on both nodes before you add SPC3 cards.

We've also updated the `show chassis high availability` information command output to include messages related to SPC3 installation.

[See [Multinode High Availability](#) and [show chassis high-availability information](#).]

J-Web

- **Enhanced Advanced Threat Prevention feature (SRX Series)**—Starting in Junos OS Release 22.2R1, we've improved the Advanced Threat Prevention page for better experience. You can configure:
 - Anti-malware profiles to define and send files to the cloud for inspection and to act when malware is detected ([Security Services > Advanced Threat Prevention > Anti-malware](#)).
 - Security Intelligence (Seclintel) profiles to work with Seclintel feeds, such as command and control (C&C), DNS, and infected hosts ([Security Services > Advanced Threat Prevention > Seclintel Profiles](#)).

- SeIntel profile groups to add SeIntel profiles and to associate the group with a security policy ([Security Services > Advanced Threat Prevention > SeIntel Profile Groups](#)).

[See [About the Anti-malware Page](#), [About the SeIntel Profiles Page](#), and [About the SeIntel Profile Groups Page](#).]

- **Support for anti-malware and SeIntel profile groups in security policy (SRX Series)**—Starting in Junos OS 22.2R1 Release, you can assign an anti-malware profile and a SeIntel profile group to security policies to inspect the profiles.

[See [Add a Rule to a Security Policy](#).]

Juniper Advanced Threat Prevention Cloud (ATP Cloud)

- **Support to configure DNS cache entries (SRX300, SRX4200, and SRX4600)**—Starting in Junos OS Release 22.2R1, you can configure a list of static benign and command-and-control (C2) domains in the DNS cache to take immediate action on configured domains.

To configure benign and C2 domains, run the commands `set services security-metadata-streaming dns-cache custom-list benign <domain>` and `set services security-metadata-streaming dns-cache custom-list c2 <domain>`

To view the benign and C2 entries in the DNS cache, use the commands `show services dns-filtering cache summary`, `show services dns-filtering cache c2`, and `show services dns-filtering cache benign`.

[See [security-metadata-streaming](#) and [show services dns-filtering cache](#).]

MPLS

- **Support for SD-WAN (SRX5000 line of devices)**—Starting in Junos OS Release 22.2R1, we support MPLS-based SD-WAN deployments on SRX5000 devices at spoke and hub locations. You can configure the SRX5000 devices to permit or deny virtual routing and forwarding (VRF)-based traffic that enters the device from overlay tunnels. The support includes:

- Networks with Layer 3 VPN
- VRF instances
- MPLS-based flow mode
- MPLS-over-GRE and MPLS-over-IPsec tunnels
- IPv4 and IPv6 traffic

[See [Configuring Security Policies for a VRF Routing Instance](#).]

Multicast

- **Support for multicast overlay on unicast SD-WAN deployments (SRX380, SRX345, SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 22.2R1, we support BGP-MVPN on existing unicast SD-WAN deployments.

[See [BGP-MVPN SD-WAN Overlay](#).]

Network Address Translation (NAT)

- **NAT support for DNS (SRX Series, vSRX, and cSRX)**—Starting in Junos OS Release 22.2R1, you can use DNS and a fully qualified domain name (FQDN) with either source NAT or destination NAT as part of your NAT configuration.

You can use DNS name servers to resolve hostnames to IP addresses. A DNS cache time to live (TTL) is introduced under the address-book option for each DNS name entry. We support a minimum DNC cache TTL of 16 seconds.

In case of multiple IP addresses in the DNS response, the first IP address in the response is added to the NAT pool.

[See [Address Books and Address Sets](#) and [show security nat source pool](#).]

Unified Threat Management (UTM)

- **Web filtering support to nonstandard ports (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 22.2R1, we've extended the Web filtering support for HTTP or HTTPS traffic to nonstandard ports.

[See [web-filtering](#) and [show security utm web-filtering status](#).]

VPNs

- **New ARI-TS routing protocol type for IPsec VPN traffic selector routes (MX-SPC3, SRX Series firewalls, and vSRX running iked process)**—Starting in Junos OS Release 22.2R1, when an IPsec negotiation is completed using a traffic selector configuration, the routes are installed as auto route insertion for traffic selectors (ARI-TS) routes instead of static routes.

Starting in Junos OS Release 22.2R1, ARI routes are considered as a routing protocol. These routes are installed with the same route preference and metric as in the previous implementation. With this approach, you can change the default route preference of the ARI-TS routes without impacting other routing protocols. You can also change the default preference value of the ARI-TS protocol per traffic selector to override the global option.

As ARI-TS is a new protocol, you may need to update routing policy statements depending on the configuration.

- To modify the default preference value with a global scope for an ARI-TS route, use the `set protocol ipsec-traffic-selector preference pref-value` command.
- To modify the preference value at each traffic selector level—that is, to configure a local preference value for an ARI-TS route, use the `set security ipsec vpn vpn-name traffic-selector ts-name preference pref-value` command.
- To add the ARI-TS protocol as the policy option along with the existing protocols such as BGP and OSFP, use the `set policy-options policy-statement policy-name term term-name from protocol ari-ts` command.

If you've configured the preference values at both global and local levels, the local preference value takes precedence.

[See [Understanding Traffic Selectors in Route-Based VPNs](#), [ipsec-traffic-selector](#), and [traffic-selector](#).]

Additional Features

Support for the following features has been extended to these platforms.

- **BGP, OSPF, and OSPFv3 authentication and encryption using manual IPsec SA** (MX240, MX480, and MX960 with MX-SPC3, SRX Series devices and vSRX running iked process). OSPF for IPv6, also known as OSPF version 3 (OSPFv3), does not have built-in authentication to ensure that routing packets are not altered and re-sent to the router. Starting in Junos OS Release 22.2R1, you can use IPsec to encrypt and secure BGP, OSPF, and OSPFv3 packets.

To configure IPsec for BGP, OSPF, and OSPFv3, define a security association (SA) with the `security-association sa-name` configuration option at the `[edit security ipsec]` hierarchy level for both MX Series and SRX Series platforms. You then apply the configured SA to the BGP, OSPF, and OSPFv3 configurations.

[See [security-association](#).]

To view the configured IPsec SAs for BGP, OSPF, and OSPFv3:

- On MX240, MX480, and MX960 with MX-SPC3, and on SRX Series devices and vSRX running the iked process, use the `show security ipsec control-plane-security-associations` command.
[See [show security ipsec control-plane-security-associations](#).]
- On MX240, MX480, and MX960 routers with MS-MPC/MS-MIC, use the `show ipsec security-associations` command.
[See [show ipsec security-associations](#).]
- On SRX Series devices running the kmd process, use the `show security ipsec security-associations` command.
[See [show security ipsec security-associations](#).]



NOTE: We do not support this feature with BGP, OSPF, and OSPFv3 over the secure tunnel (st0) interface.

[See [Understanding OSPFv3 Authentication, Using IPsec to Secure OSPFv3 Networks \(CLI Procedure\)](#), and [Example: Configuring IPsec Authentication for an OSPF Interface](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 201](#)
- [Network Address Translation \(NAT\) | 201](#)
- [Network Management and Monitoring | 202](#)
- [SSL Proxy | 202](#)
- [VPNs | 202](#)
- [VPLS | 203](#)

Learn about what changed in this release for SRX Series.

Authentication and Access Control

- **SHA-1 password format deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX and vSRX)**—We've removed the sha1 option at the [edit system login password format] hierarchy level because SHA-1 is no longer supported for plain-text password encryption.

Network Address Translation (NAT)

- **Group routing instances (SRX5600)**—Starting in Junos OS release 22.2R1, you can group the routing instances using the `routing-group` command. The `routing-group` option is added at [edit security nat destination], [edit security nat source], and [edit security nat static] hierarchies.

[See [rule-set \(Security Source NAT\)](#), [rule-set \(Security Destination NAT\)](#), and [rule-set \(Security Static NAT\)](#).]

Network Management and Monitoring

- **DES deprecation for SNMPv3 (Junos)**—The Data Encryption Standard (DES) privacy protocol for SNMPv3 is deprecated due to weak security and vulnerability to cryptographic attacks. For enhanced security, configure the triple Data Encryption Standard (3DES) or the Advanced Encryption Standard (CFB128-AES-128 Privacy Protocol) as the encryption algorithm for SNMPv3 users.

[See [privacy-3des](#) and [privacy-aes128](#).]

- **Changes to the NETCONF `<edit-config>` RPC response (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the `<edit-config>` operation returns an error, the NETCONF server does not emit a `<load-error-count>` element in the RPC response. In earlier releases, the `<edit-config>` RPC response includes the `<load-error-count>` element when the operation fails.

SSL Proxy

- **No session cache entry store during SSL session resumption (SRX Series Devices)**— When an SSL session attempts to re-initiates a full handshake and the server rejects that session resumption, the session cache does not store session information and remains empty. This issue is seen in a setup where a client device is using TLS1.1 version and the server is using TLS1.3 (maximum) version.

In Junos OS Release 22.1R1 and later releases, the session cache stores session information even when the session resumption is rejected, and you can see the session cache entries using the `show services ssl proxy session-cache entries summary` command.

VPNs

- **Deprecating IPsec Manual VPN Configuration Statement (SRX Series Devices and vSRX running kmd process)**—Starting in Junos OS Release 22.2R1, we'll be deprecating the Manual IPsec VPN (flow mode). This means that you cannot establish a manual IPsec security association (SA) using the `[edit security ipsec vpn vpn-name manual]` configuration hierarchy.

As part of this change, we'll be deprecating the `[edit security ipsec vpn vpn-name manual]` hierarchy level and its configuration options.

[See [manual](#).]

- **IPsec VPN traffic selector routes are changed from 'static routes' to 'ARI-TS' routes (MX-SPC3, SRX Series and vSRX running iked process)**—Starting in Junos OS Release 22.2R1, when an IPsec negotiation is completed using traffic selectors configuration, these routes are now installed as ARI-TS (Auto route insertion for traffic selectors) routes instead of static routes. These routes are by default installed with the same route preference and metric as the previous implementation. ARI-TS routes are inserted as '[ARI-TS/5]'.

With this approach, you can change the route preference of the ARI-TS routes without impacting other routing protocols.

[See [New ARI-TS Routing protocol](#).]

- **Include IPv6 address in a self-signed certificate (SRX Series devices and vSRX3.0)**—We support manual generation of a self-signed certificate for the given distinguished name using IPv6 address in addition to the IPv4 address that was supported earlier. Use the `request security pki local-certificate generate-self-signed` command with `ipv6-address` option to include ipv6 address in a self-signed certificate.

[See [request security pki local-certificate generate-self-signed \(Security\)](#).]

- **Unable to connect with OCSP Server for Revocation Check (SRX Series Devices and vSRX)**—When performing revocation check using OCSP, the SRX device does not attempt to connect with the OCSP server when the OCSP server URL contains a domain name that the DNS server cannot resolve. In this case, when the SRX device cannot establish connection to the OCSP server and when one of the following configuration options is set, the OCSP revocation check will either allow or fallback to using CRL:

- `set security pki ca-profile OCSP-ROOT revocation-check ocsp connection-failure disable`
- `set security pki ca-profile OCSP-ROOT revocation-check ocsp connection-failure fallback-crl`

When the SRX device cannot establish connection to the OCSP server and if these options are not configured, then the certificate validation fails.

[See [ocsp \(Security PKI\)](#).]

VPLS

- **No output byte increment on VPLS interface when configured with output filter with policer action (SRX Series Devices)**— When you upgrade your device to Junos OS Release 19.4R3-S1 or later, and the VPLS interface has an output filter with policer action applied to it, the VPLS interface does not pass the traffic. Because of this issue, the output bytes do not increment on that interface, and when

you display details using the `show interfaces <interface-name> extensive | no-more` output, the VPLS interface shows output bytes as 0. In Junos OS Release 22.2R1, the `show interfaces <interface-name> extensive | no-more` command output shows the correct details.

Known Limitations

IN THIS SECTION

- [Platform and Infrastructure | 204](#)

Learn about known limitations in Junos OS Release 22.2R1 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.

Platform and Infrastructure

- On SRX4600 device, the CPU might overrun while performing sanity check due to incompatibility issues between ukern scheduler and Linux driver which might lead to traffic loss. [PR1641517](#)

Open Issues

Learn about open issues in Junos OS Release 22.2R1 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.

Chassis Clustering

- 10GbE DAC cable is not supported at CTL or FAB link at SRX4100 and SRX4200 chassis cluster setup. [PR1636365](#)

High Availability (HA) and Resiliency

- ISSU will be aborted or failed with the warning. 'warn-message "ISSU is not supported for Clock Synchronization (SyncE)"; ''override'In '/var/tmp/paSBfY/etc/indb//config.indb' line 162included from '/var/tmp/paSBfY/etc/indb/issu.indb' line 10 'override' syntax errorISSU not supported as current image uses explicit tags for message structures. [PR1628172](#)

Interfaces and Chassis

- Traffic drop might be seen on irb interface on SRX1500 for network control forwarding class when verifying DSCP classification based on single and multiple code-points. [PR1611623](#)

Platform and Infrastructure

- With SSL proxy configured along with Web proxy, the client session might not closed on the device even though proxy session ends gracefully. [PR1580526](#)
- HA active/passive mode on-box logging in logical systems and tenant systems, the intermittently security log contents of binary log file in logical systems and tenant systems are not as expected. [PR1587360](#)
- On SRX1500 devices, ISSU is getting aborted with ISSU is not supported for Clock Synchronization (SyncE). [PR1632810](#)
- SMTPS sessions are not getting identified when traffic is sent from IXIA (BPS) profile. [PR1635929](#)
- The remote access Juniper Networks standard license might not get freed up while disconnect or reconnect after RG0 failover. [PR1642653](#)
- Firewall authentication with user firewall based RADIUS access has syslog missing the username and rule. [PR1654842](#)

Unified Threat Management (UTM)

- If only EWF is configured, there can be a performance impact due to JDPI parsing overhead. In such case, to recover the performance, Web Filter can be configured in performance mode using the set security utm default-configuration web-filtering performance-mode. [PR1653793](#)

User Interface and Configuration

- Use load update instead of load override to prevent the error messages. [PR1630315](#)

VPNs

- 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 SRX5000 line of devices, in some scenario, the device output might display obsolete IPsec SA and NHTB entry even when the peer tear down the tunnel. [PR1432925](#)
- Tunnel debugging configuration is not synchronized to the backup node. It needs to be configured again after RG0 failover. [PR1450393](#)

Resolved Issues

Learn about the issues fixed in this release for SRX Series.

Chassis Clustering

- Secondary node in a chassis cluster might go into reboot loop on SRX Series devices. [PR1606724](#)
- The Create Bearer Request might be dropped on SRX Series devices. [PR1629672](#)
- Post a series of actions MNHA functionality might not be available despite the configuration presence. [PR1638794](#)
- MSISDN prepended with additional digits in the logs. [PR1646463](#)
- Failover might not happen correctly in a chassis cluster when there is a hardware issue with the Central Point. [PR1651501](#)

Flow-Based and Packet-Based Processing

- Packets might not be classified according to the CoS rewrite configuration. [PR1634146](#)
- The process nsd might crash continuously due to failure in creating/reinitializing the file /var/db/ext/monitor-flow-cfg. [PR1638008](#)
- The traffic might get lost when using dedicated HA fabric link. [PR1651836](#)
- Performance degradation might be observed when Express Path and PME are both enabled. [PR1652025](#)

Interfaces and Chassis

- Members MAC might be different from parent reth0 interface, resulting loss of traffic. [PR1583702](#)

Intrusion Detection and Prevention (IDP)

- SRX Series devices pause when the show security idp attack attack-list policy combine-policy command is executed. [PR1616782](#)
- Packet Forwarding Engine generates core files on all Junos OS platforms. [PR1634305](#)

J-Web

- The reboot or halt from J-Web might fail on SRX series devices. [PR1638370](#)
- Significant performance improvements were made to J-Web. [PR1652676](#)

Junos XML API and Scripting

- Junos OS: Certificate validation is skipped when fetching system scripts from a HTTPS URL (CVE-2022-22156) [PR1542229](#)

Network Address Translation (NAT)

- DNS proxy service on SRX Series devices might stop working after commit operation is performed. [PR1598065](#)
- New persistent NAT or normal source NAT sessions might fail due to noncleared aged out sessions. [PR1631815](#)

Platform and Infrastructure

- CFMD core files might be seen on SRX Series devices. [PR1538173](#)
- The process pkid core files might be observed during local certificate enrollment. [PR1573892](#)
- Syslog message **%AUTH-3: warning: can't get client address: Bad file descriptor** is displayed at J-Web login. [PR1581209](#)
- BGP adjacency might not get established in Layer 2 IRB scenario. [PR1582871](#)
- Getting UNKNOWN instead of HTTP-PROXY for application and UNKNOWN instead of GOOGLE-GEN in RT-FLOW close messages These messages can be seen in the RT-flow close log and these are

due to JDPI not engaged for the session. This might affect the application identification for the web-proxy session traffic. [PR1588139](#)

- The issue is when we enable TCP path finder in the VPN gateway, VPN connection is established properly. After VPN connection is established, able to ping from JSC installed CLIENT to SERVER behind gateway, but unable to ping from SERVER behind gateway to Juniper Secure Connect installed CLIENT. [PR1611003](#)
- Execute RSI on SRX5000 line of devices with IOC2 card installed may trigger data plane failover. [PR1617103](#)
- The Layer 2 switching doesn't work as expected when running VRRP on IRB interface. [PR1622680](#)
- On SRX Series devices running DNS Security, if a DGA was detected and the action in the configuration was set to permit, under rare circumstances, a log would not be generated by the device. [PR1624076](#)
- PKID could stop and generate a core file when there was limited memory available on the Routing Engine. [PR1624613](#)
- The PKID process stops due to null pointer dereferencing during local certificate verification in some cases. [PR1624844](#)
- A major alarm DPDK Tx stuck issue of SRX4100 and SRX4200 devices. [PR1626562](#)
- Error message gencfg_cfg_msg_gen_handler drop might be seen after running commit command. [PR1629647](#)
- The srxpfe process might crash on SRX4600 device. [PR1630990](#)
- Reverse DNS lookups will no longer be stored in the DNSF cache when using DNS security. [PR1631000](#)
- The show commands to display DNS cache summary, display only DNS cache C2 entries and display only DNS cache begin entries are needed. [PR1631002](#)
- Signature package update might fail and the AppID process might stop on SRX Series devices. [PR1632205](#)
- Tasks of download manager might not be resumed post reboot. [PR1633503](#)
- On SRX Series devices running DNS Security, a dataplane memory leak may occur within the DNSF plugin when entries age-out of the DNSF cache. [PR1633519](#)
- IP monitor might install default route with incorrect preference value when multiple IP monitoring is configured. [PR1634129](#)
- Most of the Dynamic Address Entries might report 0 IPv4 entries. [PR1634881](#)

- The srxpfe process might stop while installing IDP sigpack with scaled traffic on SRX Series devices. [PR1637181](#)
- Unable to connect to domain controller on installing Microsoft KB update. [PR1637548](#)
- AppID installation failure on the secondary HA node in case of failover. [PR1638588](#)
- The spcd process might stop during certain Linux based FPC card restart. [PR1638975](#)
- The error is seen during the non ISSU upgrade from Junos OS release 15.1 to Juno OS release 18.2 and later releases. [PR1639610](#)
- Configuration change during AppQoS session might result in Packet Forwarding Engine stop with flowd process generates core file. [PR1640768](#)
- Traffic might be dropped due to the RX queue being full. [PR1641793](#)
- Observing Error usp_ipc_client_recv_::ipc_pipe_read() due to core file, when checking show security monitoring CLI command. [PR1641995](#)
- The Packet Forwarding Engine process might pause on SRX Series devices. [PR1642914](#)
- The ATP integrated service might get impacted on SRX Series devices with logical system. [PR1643373](#)
- The on-box security logs might be not storing the session-id as a 64-bit integer, resulting in incorrect session-id's being present in the on-box logs. [PR1644867](#)
- Issue with the command clear security idp counters packet-log logical-system all. [PR1648187](#)
- The severity of AAMW and SMS control and submission channel alarms have been reduced from major to minor to avoid triggering a chassis cluster failover in the event of an upstream network issue. [PR1648330](#)
- SCB reset with Error : zfchip_scan line = 844 name = failed due to PIO errors. [PR1648850](#)
- Unable to get the firewall-authentication users details on node 1. [PR1651129](#)
- SMB file submissions to ATP cloud failed. [PR1653098](#)
- Certificate based VPN tunnel is not established. [PR1655571](#)
- Radius responses that take longer than 15 seconds can cause SRX Series devices to declare authentication failure. [PR1658833](#)

Routing Protocols

- Delay in BGP session establishment due to longer time for the listening task to be ready on all platforms running rpd. [PR1651211](#)

Unified Threat Management (UTM)

- New UTM content filtering CLI is changing from seclog to log. [PR1634580](#)
- Modification of content filtering rule order after Junos OS release 21.4 would not have the desired effect. [PR1653488](#)

User Interface and Configuration

- In an SRX Series devices with chassis cluster and VPN configuration, primary node in cluster might generate kmd core files in a loop when a commit fails with lock can not be taken on other node followed by another commit. [PR1608718](#)
- MGD core might be observed upon ISSU upgrade. [PR1632853](#)
- Unable to access configure exclusive mode after mgd process is stopped. [PR1641025](#)

VPNs

- The configuration change in SRG-1 might cause HA link encryption tunnel flap. [PR1598338](#)
- The process iked stop might be seen for IKEv1 based VPN tunnels. [PR1608724](#)
- Fragmented packets might drop when PMI is enabled. [PR1624877](#)
- Traffic loss over IPsec tunnel might be seen on SRX Series devices. [PR1628007](#)
- IPsec tunnel might stop processing traffic. [PR1636458](#)
- The kmd process might crash if the IKE negotiation fragment packets are missed during initiating an IKE SA rekey. [PR1638437](#)
- IPsec tunnel through IPv6 won't establish after rebooting. [PR1653704](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 211](#)

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

Starting in Junos OS release 21.2R1, all Junos OS products which were previously running on FreeBSD 11.x based Junos OS are migrated to FreeBSD 12.x based Junos OS, except EX4400. Starting with Junos OS release 21.3R1, EX4400 platforms are migrated to FreeBSD 12.x based Junos OS.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

For information about ISSU, see the [Chassis Cluster User Guide for Security Devices](#).

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

We have two types of releases, EOL and EEOL:

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

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

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 12: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Junos OS Release Notes for vMX

IN THIS SECTION

- [What's New | 213](#)
- [What's Changed | 215](#)
- [Known Limitations | 217](#)
- [Open Issues | 218](#)
- [Resolved Issues | 218](#)
- [Upgrade Instructions | 218](#)

These release notes accompany Junos OS Release 22.2R1 for vMX. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [EVPN | 213](#)
- [Junos Telemetry Interface | 213](#)
- [MPLS | 213](#)
- [OpenConfig | 214](#)
- [Routing Protocols | 214](#)

Learn about new features introduced in this release for vMX.

EVPN

- **EVPN active/active redundancy, aliasing, and mass MAC withdrawal (MX Series and vMX)**—Starting in Junos OS Release 22.2R1, the listed devices support EVPN active/active redundancy, aliasing, and mass MAC withdrawal, integrated with VXLAN in the data plane. These features provide resilient inter-data center connectivity to the established Data Center Interconnect (DCI) technologies. This new support builds an end-to-end DCI solution by integrating EVPN active/active multicast with DP VXLAN.

Use existing configuration statements to configure active/active redundancy at the ESI level on the loopback (lo0) interface. Include lo0 as the virtual tunnel endpoint (VTEP) interface in the routing instance.

[See [EVPN-over-VXLAN Supported Functionality](#).]

Junos Telemetry Interface

- **Support for QoS sensor (MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, and vMX)**—Junos OS Release 22.2R1 introduces QoS sensors for JTI based on the OpenConfig data model **open-config-qos version 0.3.0**.

[See [Telemetry Sensor Explorer](#).]

MPLS

- **Support for ingress and transit chained CNHs for BGP Labeled Unicast (BGP-LU) IPv4 (MX204, MX480, MX960, MX10003, and vMX)**—Starting in Junos OS Release 22.2R1, you can configure the

chained composite next hops (CNHs) for devices handling ingress or transit traffic in the network. We've added support only for the following options on the listed MX Series devices:

- BGP-LU for IPv4 on the ingress router—set routing-options forwarding-table chained-composite-next-hop ingress labeled-bgp inet
- BGP-LU on the transit router—set routing-options forwarding-table chained-composite-next-hop transit labeled-bgp

You can also configure class of service (CoS) and define rewrite rules for ingress and transit chained CNHs for BGP-LU.



NOTE: This feature is supported only on MPC9 and the previous models of line cards.

OpenConfig

- **Support for OpenConfig telemetry system configuration (ACX5448, ACX710, MX204, MX240, MX150, MX960, MX10004, MX2008, MX2010, MX2020, PTX10002, and vMX)—Junos OS Release 22.2R1** introduces support for OpenConfig data models **openconfig-telemetry.yang** and **openconfig-telemetry-types.yang**. The support includes streaming of state data for dynamic subscriptions and OpenConfig configuration for persistent subscriptions.

[See [Mapping OpenConfig Telemetry System Model Commands to Junos Configuration](#).]

Routing Protocols

- **DCSPF support for SR-TE with Flex Algo (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, vMX, PTX1000, PTX3000, and PTX5000)—Starting in Junos OS Release 22.2R1, we support the flexible algorithm (Flex Algo) as a constraint in the compute profile of a segment routing-traffic engineering (SR-TE) LSP. The computation combines any constraints in the compute profile with the ones in the Flex Algo definition to find the resultant path. It uses the Flex Algo segment identifiers (SIDs) in the configuration to compress the resultant path.**

We support the feature only for IPv4 SR-MPLS SIDs. You can use SR-TE policy constraints to further fine-tune Flex Algo constraints.

[See [Enabling Distributed CSPF for Segment Routing LSPs](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 215](#)
- [General Routing | 215](#)
- [Network Management and Monitoring | 216](#)
- [User Interface and Configuration | 216](#)

Learn about what changed in this release for vMX.

Authentication and Access Control

- **SHA-1 password format deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX and vSRX)**—We've removed the sha1 option at the [edit system login password format] hierarchy level because SHA-1 is no longer supported for plain-text password encryption.

General Routing

- **The <request-system-zeroize/> RPC response indicates when the device successfully initiates the requested operation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the <request-system-zeroize/> RPC successfully initiates the zeroize operation, the device emits the <system-zeroize-status>zeroizing re0</system-zeroize-status> response tag to indicate that the process has started. If the device fails to initiate the zeroize operation, the device does not emit the <system-zeroize-status> response tag.
- **Instance type change is not permitted from default to L3VRF in open configuration (ACX Series, EX Series, MX Series, QFX Series, SRX Series, vMX, and vSRX)**—DEFAULT_INSTANCE is the primary instance that runs when there is no specific instance type configured in the route set routing-options. Any instance you explicitly configure is translated into set routing-instance r1 routing-options. The issue appears in translation, when you change instance type DEFAULT_INSTANCE (any instance to DEFAULT_INSTANCE) to L3VRF or L3VRF to DEFAULT_INSTANCE. As a result, such changes are not permitted. Additionally, DEFAULT_INSTANCE can only be named DEFAULT, and DEFAULT is reserved for DEFAULT_INSTANCE, therefore allowing no such changes.

Network Management and Monitoring

- **DES deprecation for SNMPv3**—The Data Encryption Standard (DES) privacy protocol for SNMPv3 is deprecated due to weak security and vulnerability to cryptographic attacks. For enhanced security, configure the triple Data Encryption Standard (3DES) or the Advanced Encryption Standard (CFB128-AES-128 Privacy Protocol) as the encryption algorithm for SNMPv3 users.

[See [privacy-3des](#) and [privacy-aes128](#).]

- **Changes when deactivating or deleting instances of the ephemeral configuration database (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The following changes apply when you deactivate or delete ephemeral database instances in the static configuration database:

- When you deactivate the entire [edit system configuration-database ephemeral] hierarchy level, the device deletes the files and corresponding configuration data for all user-defined ephemeral instances. In earlier releases, the files and configuration data are preserved; however, the configuration data is not merged with the static configuration database.
- When you delete an ephemeral instance in the static configuration database, the instance's configuration files are also deleted. In earlier releases, the configuration files are preserved.
- You can delete the files and corresponding configuration data for the default ephemeral database instance by configuring the delete-ephemeral-default statement in conjunction with the ignore-ephemeral-default statement at the [edit system configuration-database ephemeral] hierarchy level.

[See [Enable and Configure Instances of the Ephemeral Configuration Database](#).]

- **Changes to the NETCONF <edit-config> RPC response (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the <edit-config> operation returns an error, the NETCONF server does not emit a <load-error-count> element in the RPC response. In earlier releases, the <edit-config> RPC response includes the <load-error-count> element when the operation fails.

User Interface and Configuration

- **Load JSON configuration data with unordered list entries (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—The Junos schema requires that list keys precede any other siblings within a list entry and appear in the order specified by the schema. Junos devices provide two options to load JSON configuration data that contains unordered list entries:
 - Use the request system convert-json-configuration operational mode command to produce JSON configuration data with ordered list entries before loading the data on the device.

- Configure the `reorder-list-keys` statement at the `[edit system configuration input format json]` hierarchy level. After you configure the statement, you can load JSON configuration data with unordered list entries, and the device reorders the list keys as required by the Junos schema during the load operation.
- When you configure the `reorder-list-keys` statement, the load operation can take significantly longer to parse the configuration, depending on the size of the configuration and number of lists. Therefore, for large configurations or configurations with many lists, we recommend using the `request system convert-json-configuration` command instead of the `reorder-list-keys` statement.

[See [json](#) and [request system convert-json-configuration](#)]

- Junos XML protocol Perl modules deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—We no longer provide the Junos XML protocol Perl client for download. To use Perl to manage Junos devices, use the NETCONF Perl library instead.

[See [Understanding the NETCONF Perl Client and Sample Scripts..](#)]

Known Limitations

IN THIS SECTION

- Platform and Infrastructure | [217](#)

Learn about known limitations in Junos OS Release 22.2R1 for vMX.

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

Platform and Infrastructure

- The riot might crash due to a rare issue if vMX run in the performance mode. [PR1534145](#)

Open Issues

There are no known issues in hardware and software in Junos OS Release 22.2R1 for vMX.

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

Resolved Issues

There are no resolved issues in Junos OS Release 22.2R1 for vMX.

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.

Starting in Junos OS release 21.2R1, all Junos OS products which were previously running on FreeBSD 11.x based Junos OS are migrated to FreeBSD 12.x based Junos OS, except EX4400. Starting with Junos OS release 21.3R1, EX4400 platforms are migrated to FreeBSD 12.x based Junos OS.

Junos OS Release Notes for vRR

IN THIS SECTION

- [What's New | 219](#)
- [What's Changed | 219](#)
- [Known Limitations | 220](#)
- [Open Issues | 220](#)
- [Resolved Issues | 220](#)

These release notes accompany Junos OS Release 22.2R1 for vRR. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [Routing Protocols | 219](#)

Learn about new features introduced in this release for vRR.

Routing Protocols

- **TCP-AO for RPKI validation sessions (MX204, MX240, MX480, MX960, MX10003, MX10008, MX10016, MX2008, MX2010, MX2020, PTX1000, PTX10002, PTX10008, PTX10016, and vRR) —** Starting in Junos OS Release 22.2R1, you can use TCP Authentication Option (TCP-AO) to authenticate resource public key infrastructure (RPKI) validation sessions for securing the Internet's routing infrastructure, such as BGP. Using RPKI, legitimate holders of Internet number resources can control the operation of Internet routing protocols to prevent route hijacking and other attacks.

To enable a TCP-AO chain to authenticate an RPKI validation session, use `authentication-algorithm ao` and the configured `authentication-key-chain keychain` at the `[edit routing-options validation group group_name session address]` and `[edit routing-options validation group group_name]` hierarchy levels.

See [\[TCP Authentication Option \(TCP-AO\).\]](#)

What's Changed

There are no changes in behavior and syntax in Junos OS Release 22.2R1 for vRR.

Known Limitations

There are no known limitations in hardware and software in Junos OS 22.2R1 for vRR.

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

To learn more about common BGP or routing known limitations in Junos OS 22.2R1, see ["Known Limitations" on page 100](#) for MX Series routers.

Open Issues

There are no known issues in hardware and software in Junos OS Release 22.2R1 for vRR.

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

Resolved Issues

IN THIS SECTION

- [Platform and Infrastructure | 220](#)

Learn about the issues fixed in this release for vRR.

Platform and Infrastructure

- The video console for vRR might not work after an upgrade to Junos with upgraded FreeBSD.
[PR1644806](#)

Junos OS Release Notes for vSRX

IN THIS SECTION

- [What's New | 221](#)
- [What's Changed | 223](#)
- [Known Limitations | 225](#)
- [Open Issues | 225](#)
- [Resolved Issues | 226](#)
- [Migration, Upgrade, and Downgrade Instructions | 228](#)

These release notes accompany Junos OS Release 22.2R1 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

- [Flow-Based and Packet-Based Processing | 222](#)
- [Network Address Translation \(NAT\) | 222](#)
- [Unified Threat Management \(UTM\) | 222](#)
- [VPNs | 222](#)

Learn about new features introduced in this release for vSRX.

Flow-Based and Packet-Based Processing

- **Support for IPv6 tunnel (SRX Series and vSRX 3.0)**— Starting in Junos OS Release 22.2R1, you can encapsulate IPv4 and IPv6 traffic over the IPv6 network.

The IPv6 tunnel helps IPv4 traffic traverse over the IPv6 network. You can use IPv6 tunneling in various features such as policy routing and preferential billing. For example, a set-top box that supports only IPv4 traffic can traverse the server over an IPv6 network.

[See [show security flow session](#).]

Network Address Translation (NAT)

- **NAT support for DNS (SRX Series, vSRX, and cSRX)**—Starting in Junos OS Release 22.2R1, you can use DNS and a fully qualified domain name (FQDN) with either source NAT or destination NAT as part of your NAT configuration.

You can use DNS name servers to resolve hostnames to IP addresses. A DNS cache time to live (TTL) is introduced under the address-book option for each DNS name entry. We support a minimum DNC cache TTL of 16 seconds.

In case of multiple IP addresses in the DNS response, the first IP address in the response is added to the NAT pool.

[See [Address Books and Address Sets](#) and [show security nat source pool](#).]

Unified Threat Management (UTM)

- **Web filtering support to nonstandard ports (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 22.2R1, we've extended the Web filtering support for HTTP or HTTPS traffic to nonstandard ports.

[See [web-filtering](#) and [show security utm web-filtering status](#).]

VPNs

- **New ARI-TS routing protocol type for IPsec VPN traffic selector routes (MX-SPC3, SRX Series firewalls, and vSRX running iked process)**—Starting in Junos OS Release 22.2R1, when an IPsec negotiation is completed using a traffic selector configuration, the routes are installed as auto route insertion for traffic selectors (ARI-TS) routes instead of static routes.

Starting in Junos OS Release 22.2R1, ARI routes are considered as a routing protocol. These routes are installed with the same route preference and metric as in the previous implementation. With this approach, you can change the default route preference of the ARI-TS routes without impacting other routing protocols. You can also change the default preference value of the ARI-TS protocol per traffic selector to override the global option.

As ARI-TS is a new protocol, you may need to update routing policy statements depending on the configuration.

- To modify the default preference value with a global scope for an ARI-TS route, use the `set protocol ipsec-traffic-selector preference pref-value` command.
- To modify the preference value at each traffic selector level—that is, to configure a local preference value for an ARI-TS route, use the `set security ipsec vpn vpn-name traffic-selector ts-name preference pref-value` command.
- To add the ARI-TS protocol as the policy option along with the existing protocols such as BGP and OSFP, use the `set policy-options policy-statement policy-name term term-name from protocol ari-ts` command.

If you've configured the preference values at both global and local levels, the local preference value takes precedence.

[See [Understanding Traffic Selectors in Route-Based VPNs](#), [ipsec-traffic-selector](#), and [traffic-selector](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 223](#)
- [Network Management and Monitoring | 224](#)
- [VPNs | 224](#)

Learn about what changed in this release for vSRX.

Authentication and Access Control

- **SHA-1 password format deprecated (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX and vSRX)**—We've removed the `sha1` option at the `[edit system login password format]` hierarchy level because SHA-1 is no longer supported for plain-text password encryption.

Network Management and Monitoring

- **Changes to the NETCONF <edit-config> RPC response (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When the <edit-config> operation returns an error, the NETCONF server does not emit a <load-error-count> element in the RPC response. In earlier releases, the <edit-config> RPC response includes the <load-error-count> element when the operation fails.

VPNs

- **Deprecating IPsec Manual VPN Configuration Statement (SRX Series Devices and vSRX running kmd process)**—Starting in Junos OS Release 22.3R1, we'll be deprecating the Manual IPsec VPN (flow mode). This means that you cannot establish a manual IPsec security association (SA) using the [edit security ipsec vpn *vpn-name* manual] configuration hierarchy.

As part of this change, we'll be deprecating the [edit security ipsec vpn *vpn-name* manual] hierarchy level and its configuration options.

[See [manual](#).]

- **IPsec VPN traffic selector routes are changed from 'static routes' to 'ARI-TS' routes (MX-SPC3, SRX Series and vSRX running iked process)**—Starting in Junos OS Release 22.2R1, when an IPsec negotiation is completed using traffic selectors configuration, these routes are now installed as ARI-TS (Auto route insertion for traffic selectors) routes instead of static routes. These routes are by default installed with the same route preference and metric as the previous implementation. ARI-TS routes are inserted as '[ARI-TS/5]'.

With this approach, you can change the route preference of the ARI-TS routes without impacting other routing protocols.

[See [New ARI-TS Routing protocol](#).]

- **Include IPv6 address in a self-signed certificate (SRX Series devices and vSRX3.0)**—We support manual generation of a self-signed certificate for the given distinguished name using IPv6 address in addition to the IPv4 address that was supported earlier. Use the request security pki local-certificate generate-self-signed command with ipv6-address option to include ipv6 address in a self-signed certificate.

[See [request security pki local-certificate generate-self-signed \(Security\)](#).]

- **Unable to connect with OCSP Server for Revocation Check (SRX Series Devices and vSRX)**—When performing revocation check using OCSP, the SRX device does not attempt to connect with the OCSP server when the OCSP server URL contains a domain name that the DNS server cannot resolve. In this case, when the SRX device cannot establish connection to the OCSP server and when

one of the following configuration options is set, the OCSP revocation check will either allow or fallback to using CRL:

- `set security pki ca-profile OCSP-ROOT revocation-check ocsp connection-failure disable`
- `set security pki ca-profile OCSP-ROOT revocation-check ocsp connection-failure fallback-crl`

When the SRX device cannot establish connection to the OCSP server and if these options are not configured, then the certificate validation fails.

[See [ocsp \(Security PKI\)](#).]

Known Limitations

There are no known limitations in hardware and software in Junos OS 22.2R1 for vSRX.

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

Open Issues

Learn about open issues in Junos OS Release 22.2R1 for vSRX.

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

- The ICMPv6 TCP sequence information is missing in the ICMPv6 error generated. [PR1611202](#)

Platform and Infrastructure

- With SSL proxy configured along with Web proxy, the client session might not close on the device even though proxy session ends gracefully. [PR1580526](#)
- AMR when it is enabled in non-cso v6 over v6 mode with IPsec tunnels, the first session after reboot or forward restart, will not have multipath treatment, post that the feature works fine. [PR1643570](#)

VPNs

- In certain cases, the PUSH ACK message from the group member to the group key server might be lost. The group member can still send rekey requests for the TEK SAs before the hard lifetime expiry. Only if the key server sends any new PUSH messages to the group members, those updates would not be received by the group member since the key server would have removed the member from registered members list. [PR1608290](#)

Resolved Issues

Learn about the issues fixed in this release for vSRX.

Flow-Based and Packet-Based Processing

- Traffic in the power mode still passthrough when the ingress logic interface is manually disabled. [PR1604144](#)

Intrusion Detection and Prevention (IDP)

- SRX Series devices pause when the show security idp attack attack-list policy combine-policy command is executed. [PR1616782](#)

J-Web

- J-Web might only allow certain types of interfaces to be added in a routing instance. [PR1637917](#)
- Significant performance improvements were made to J-Web. [PR1652676](#)

Platform and Infrastructure

- Tag "RT_FLOW_SESSION_XXX" is missing in stream mode. [PR1565153](#)
- PKID core during auto-re-enrollment of CMPv2 certificates. [PR1580442](#)
- Getting UNKNOWN instead of HTTP-PROXY for application and UNKNOWN instead of GOOGLE-GEN in RT-FLOW close messages. These messages can be seen in the RT-flow close log and these are due to JDPI not engaged for the session. This may affect the app identification for the web-proxy session traffic. [PR1588139](#)
- During SaaS probing, due to race condition between APP entry addition and session processing, this core is seen. [PR1622787](#)

- On SRX Series devices running DNS security, if a DGA was detected and the action in the configuration was set to permit, under rare circumstances, a log would not be generated by the device. [PR1624076](#)
- Signature package update might fail and the AppID process might stop on SRX Series devices. [PR1632205](#)
- On SRX Series devices running DNS Security, a dataplane memory leak might occur within the DNSF plugin when entries age-out of the DNSF cache. [PR1633519](#)
- Application group name is not found for micro applications in CLI show output. [PR1640040](#)
- The junos-ssl-term is not found in ssl-trace-new logs. [PR1640075](#)
- The Packet Forwarding Engine process might pause on SRX Series devices. [PR1642914](#)
- Certificate based VPN tunnel is not established. [PR1655571](#)

Subscriber Access Management

- Same set of ciphers used in all 3 cipher categories low or medium or strong. [PR1646260](#)

Unified Threat Management (UTM)

- New UTM content filtering CLI is changing from seclog to log. [PR1634580](#)
- Modification of content filtering rule order after Junos OS release 21.4 would not have the desired effect. [PR1653488](#)
- Web browser traffic might get blocked when matched to the content filtering rule with file-types 7z. [PR1656266](#)

User Interface and Configuration

- In an SRX Series devices with chassis cluster and VPN configuration, primary node in cluster might generate kmd core files in a loop when a commit fails with lock can not be taken on other node followed by another commit. [PR1608718](#)

VPNs

- Unable to set DynamoDB in HSM module. [PR1599069](#)
- IPsec tunnel might stop processing traffic. [PR1636458](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 234](#)

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 22.2R1 for vSRX using J-Web (see [J-Web](#)) or the Junos Space Network Management Platform (see [Junos Space](#)).

Starting in Junos OS release 21.2R1, all Junos OS products which were previously running on FreeBSD 11.x based Junos OS are migrated to FreeBSD 12.x based Junos OS, except EX4400. Starting with Junos OS release 21.3R1, EX4400 platforms are migrated to FreeBSD 12.x based Junos OS.

Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Releases 17.4, 18.1, 18.2, 18.3, 18.4, 19.1, 19.2 and 19.4 is supported.

The following limitations apply:

- Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Release 19.3 and higher is not supported. For upgrade between other combinations of Junos OS Releases in vSRX and vSRX 3.0, the general Junos OS upgrade policy applies.
- The file system mounted on /var usage must be below 14% of capacity.

Check this using the following command:

```
show system storage | match " /var$" /dev/vtbd1s1f
2.7G      82M      2.4G      3%  /var
```

Using the `request system storage cleanup` command might help reach that percentage.

- The Junos OS upgrade image must be placed in the directory `/var/host-mnt/var/tmp/`. Use the `request system software add /var/host-mnt/var/tmp/<upgrade_image>`
- We recommend that you deploy a new vSRX virtual machine (VM) instead of performing a Junos OS upgrade. That also gives you the option to move from vSRX to the newer and more recommended vSRX 3.0.

- Ensure to back up valuable items such as configurations, license-keys, certificates, and other files that you would like to keep.



NOTE: For ESXi deployments, the firmware upgrade from Junos OS Release 15.1X49-Dxx to Junos OS releases 17.x, 18.x, or 19.x is not recommended if there are more than three network adapters on the 15.1X49-Dxx vSRX instance. If there are more than three network adapters and you want to upgrade, then we recommend that you either delete all the additional network adapters and add the network adapters after the upgrade or deploy a new vSRX instance on the targeted OS version.

Upgrading Software Packages

To upgrade the software using the CLI:

1. Download the **Junos OS Release 22.2R1 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 22.2R1 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-x86-64-22.2-2022-10-12.0_RELEASE_22.2_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-x86-64-22.2-2022-10-12.0_RELEASE_22.2_THROTTLE.tgz no-copy no-validate reboot
Verified junos-vsrx-x86-64-22.2-2022-10-12.0_RELEASE_22.2_THROTTLE signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING: This package will load JUNOS 22.2 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-22.2-2022-10-12.0_RELEASE_22.2_THROTTLE.tgz
Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31 junos-srx-mr-vsrx-22.2-2022-10-12.0_RELEASE_22.2_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31 junos-srx-mr-vsrx-22.2-2022-10-12.0_RELEASE_22.2_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-22.2-2022-10-12.0_RELEASE_22.2_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-22.2-2022-10-12.0_RELEASE_22.2_THROTTLE-linux.tgz ...
upgrade_platform: Input package /var/tmp/junos-srx-mr-
vsrx-22.2-2022-10-12.0_RELEASE_22.2_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-22.2-2022-10-12.0_RELEASE_22.2_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-22.2-2022-10-12.0_RELEASE_22.2_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 22.2R1 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 22.2-2022-10-12.0_RELEASE_22.2_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: 22.2-2022-10-12.0_RELEASE_22.2_THROTTLE
JUNOS OS Kernel 64-bit [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs [20171012.170745_fbsd-builder_stable_11]
JUNOS OS runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS OS time zone information [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs compat32 [20171012.170745_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20171012.170745_fbsd-builder_stable_11]
JUNOS py extensions [20171017.110007_ssd-builder_release_174_throttle]
JUNOS py base [20171017.110007_ssd-builder_release_174_throttle]
JUNOS OS vmguest [20171012.170745_fbsd-builder_stable_11]
JUNOS OS crypto [20171012.170745_fbsd-builder_stable_11]
JUNOS network stack and utilities [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Web Management Platform Package [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS common platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS mtx network modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp modules [20171017.110007_ssd-builder_release_174_throttle]
```

```
JUNOS srxtvp libs [20171017.110007_ssd-builder_release_174_throttle]  
JUNOS srx libs [20171017.110007_ssd-builder_release_174_throttle]  
JUNOS srx Data Plane Crypto Support [20171017.110007_ssd-builder_release_174_throttle]  
JUNOS daemons [20171017.110007_ssd-builder_release_174_throttle]  
JUNOS srx daemons [20171017.110007_ssd-builder_release_174_throttle]  
JUNOS Online Documentation [20171017.110007_ssd-builder_release_174_throttle]  
JUNOS jail runtime [20171012.170745_fbsd-builder_stable_11]  
JUNOS FIPS mode utilities [20171017.110007_ssd-builder_release_174_throttle]
```

Validating the OVA Image

If you have downloaded a vSRX .ova image and need to validate it, see [Validating the vSRX .ova File for VMware](#).

Note that only .ova (VMware platform) vSRX images can be validated. The .qcow2 vSRX images for use with KVM cannot be validated the same way. File checksums for all software images are, however, available on the download page.

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

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

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

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases.

Table 13: EOL and EEOL Releases

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

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Licensing

In 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that Juniper Networks has developed over the past several years.

The major components of the framework are:

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

For information about the list of supported products, see [Juniper Flex Program](#).

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you to explore software feature information to find the right software release and product for your network.
<https://apps.juniper.net/feature-explorer/>
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved.
<https://prsearch.juniper.net/InfoCenter/index?page=prsearch>
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms.
<https://apps.juniper.net/hct/home>



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

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#).
<https://pathfinder.juniper.net/compliance/>

Requesting Technical Support

IN THIS SECTION

- [Self-Help Online Tools and Resources | 237](#)
- [Creating a Service Request with JTAC | 238](#)

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/content/dam/www/assets/resource-guides/us/en/jtac-user-guide.pdf>.
- Product warranties—For product warranty information, visit <https://support.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://support.juniper.net/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://supportportal.juniper.net/s/knowledge>
- Download the latest versions of software and review release notes: <https://support.juniper.net/support/downloads/>
- Search technical bulletins for relevant hardware and software notifications: <https://supportportal.juniper.net/s/knowledge>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>

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 [Juniper Support Portal: Case Management, Product Support & More](#).
- 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

9 January 2026—Revision 24, Junos OS Release 22.2R1.

15 August 2025—Revision 23, Junos OS Release 22.2R1.

29 May 2025—Revision 22, Junos OS Release 22.2R1.

3 April 2025—Revision 21, Junos OS Release 22.2R1.

27 March 2025—Revision 20, Junos OS Release 22.2R1.

3 June 2024—Revision 19, Junos OS Release 22.2R1.

15 February 2024—Revision 18, Junos OS Release 22.2R1.

18 January 2024—Revision 17, Junos OS Release 22.2R1.

12 December 2023—Revision 16, Junos OS Release 22.2R1.

20 July 2023—Revision 15, Junos OS Release 22.2R1.

1 June 2023—Revision 14, Junos OS Release 22.2R1.

4 May 2023—Revision 13, Junos OS Release 22.2R1.

30 March 2023—Revision 12, Junos OS Release 22.2R1.

9 March 2023—Revision 11, Junos OS Release 22.2R1.

16 February 2023—Revision 10, Junos OS Release 22.2R1.

25 November 2022—Revision 9, Junos OS Release 22.2R1.

17 November 2022—Revision 8, Junos OS Release 22.2R1.

7 November 2022—Revision 7, Junos OS Release 22.2R1.

2 November 2022—Revision 6, Junos OS Release 22.2R1.

29 September 2022—Revision 5, Junos OS Release 22.2R1.

11 August 2022—Revision 3, Junos OS Release 22.2R1.

30 June 2022—Revision 2, Junos OS Release 22.2R1.

23 June 2022—Revision 1, Junos OS Release 22.2R1.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2026 Juniper Networks, Inc. All rights reserved.