

Release Notes

Published
2025-09-18

Junos OS Release 24.2R1®

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, cPCE, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, MX Series, NFX Series, QFX Series, SRX Series Firewalls, and vSRX Virtual Firewall. This release notes accompany Junos OS Release 24.2R1. 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

Introduction | 1

Junos OS Release Notes for ACX Series

What's New | 1

Junos Telemetry Interface | 2

Network Management and Monitoring | 2

Precision Time Protocol (PTP) | 3

Software Installation and Upgrade | 3

Additional Features | 3

What's Changed | 3

Known Limitations | 8

Open Issues | 8

Resolved Issues | 9

Migration, Upgrade, and Downgrade Instructions | 10

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 11

Junos OS Release Notes for cPCE

What's New | 12

Additional Features | 12

What's Changed | 13

Known Limitations | 13

Open Issues | 13

Resolved Issues | 13

Junos OS Release Notes for cRPD

What's New | 14

| Routing Protocols | 14

Known Limitations | 14

Open Issues | 15

Resolved Issues | 15

Junos OS Release Notes for cSRX

What's New | 16

| Network Management and Monitoring | 16

What's Changed | 17

Known Limitations | 17

Open Issues | 17

Resolved Issues | 17

Junos OS Release Notes for EX Series

What's New 24.2R1-S1 | 18

| Hardware | 18

What's New | 19

| Authentication and Access Control | 20

| EVPN | 20

| Interfaces | 27

| Junos OS API and Scripting | 27

| Junos Telemetry Interface | 27

| Network Management and Monitoring | 28

| Routing Policy and Firewall Filters | 29

| Routing Protocols | 29

| Serviceability | 31

| Software Installation and Upgrade | 31

| System Logging | 32

Virtual Chassis | 32

Additional Features | 32

What's Changed | 34

Known Limitations | 37

Open Issues | 38

Resolved Issues | 41

Migration, Upgrade, and Downgrade Instructions | 43

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life
Releases | 43

Junos OS Release Notes for JRR Series

What's New | 44

What's Changed | 45

Known Limitations | 45

Open Issues | 45

Resolved Issues | 46

Migration, Upgrade, and Downgrade Instructions | 46

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life
Releases | 46

Junos OS Release Notes for MX Series

What's New | 48

Hardware | 49

Authentication and Access Control | 65

Chassis | 65

Class of Service | 66

Dynamic Host Configuration Protocol | 67

EVPN | 67

Forwarding Options | 69

High Availability | 69

Interfaces | 69

Juniper Extension Toolkit (JET) | 70

Junos OS API and Scripting | 70

Junos Telemetry Interface | 71

MPLS | 75

Network Management and Monitoring | 80

Precision Time Protocol (PTP) | 83

Public Key Infrastructure (PKI) | 84

Routing Policy and Firewall Filters | 84

Routing Protocols | 85

Securing GTP and SCTP Traffic | 89

Serviceability | 89

Services Applications | 90

Source Packet Routing in Networking (SPRING) or Segment Routing | 91

Software Installation and Upgrade | 92

Subscriber Management and Services | 94

Additional Features | 94

What's Changed | 95

Known Limitations | 101

Open Issues | 103

Resolved Issues | 111

Migration, Upgrade, and Downgrade Instructions | 121

Junos OS Release Notes for NFX Series

What's New | 127

Dynamic Host Configuration Protocol | 127

Platform and Infrastructure	127
Software Installation and Upgrade	128
VPNs	128

What's Changed | 128

Known Limitations | 129

Open Issues | 129

Resolved Issues | 130

Migration, Upgrade, and Downgrade Instructions | 131

Junos OS Release Notes for QFX Series

What's New | 134

Chassis	135
EVPN	135
Junos OS API and Scripting	139
Junos Telemetry Interface	139
MPLS	140
Network Management and Monitoring	141
Routing Protocols	141
Software Installation and Upgrade	142
Additional Features	143

What's Changed | 144

Known Limitations | 148

Open Issues | 148

Resolved Issues | 149

Migration, Upgrade, and Downgrade Instructions | 151

Junos OS Release Notes for Juniper Secure Connect

What's New | 165

What's Changed | 165

Known Limitations | 165

Open Issues | 165

Resolved Issues | 166

Junos OS Release Notes for SRX Series Firewalls

What's New | 166

Hardware | 168

Application Identification (AppID) | 178

Chassis | 180

Flow-Based and Packet-Based Processing | 180

High Availability | 181

Interfaces | 181

Juniper Advanced Threat Prevention Cloud (ATP Cloud) | 182

Juniper Extension Toolkit (JET) | 182

J-Web | 182

Network Management and Monitoring | 183

Public Key Infrastructure (PKI) | 183

Serviceability | 184

Software Installation and Upgrade | 184

VPNs | 185

Additional Features | 186

What's Changed | 187

Known Limitations | 192

Open Issues | 193

Resolved Issues | 195

Migration, Upgrade, and Downgrade Instructions | 197

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 198

Documentation Updates | 199**Junos OS Release Notes for vSRX****What's New | 199**

Application Identification (AppID) | 200

High Availability | 201

Juniper Advanced Threat Prevention Cloud (ATP Cloud) | 202

Network Management and Monitoring | 202

Platform and Infrastructure | 203

VPNs | 203

What's Changed | 204**Known Limitations | 207****Open Issues | 207****Resolved Issues | 208****Migration, Upgrade, and Downgrade Instructions | 208**

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 215

Licensing | 216**Finding More Information | 216****Requesting Technical Support | 217****Revision History | 218**

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, cPCE, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, MX Series, NFX Series, QFX Series, SRX Series Firewall, and vSRX Virtual Firewall. This release notes accompany Junos OS Release 24.2R1. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [What's New | 1](#)
- [What's Changed | 3](#)
- [Known Limitations | 8](#)
- [Open Issues | 8](#)
- [Resolved Issues | 9](#)
- [Migration, Upgrade, and Downgrade Instructions | 10](#)

What's New

IN THIS SECTION

- [Junos Telemetry Interface | 2](#)
- [Network Management and Monitoring | 2](#)
- [Precision Time Protocol \(PTP\) | 3](#)
- [Software Installation and Upgrade | 3](#)
- [Additional Features | 3](#)

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

To view features supported on the ACX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Release 24.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- [ACX710](#)
- [ACX5448-D](#)
- [ACX5448-M](#)
- [ACX5448](#)

Junos Telemetry Interface

- OpenConfig MAC address and MAC address and IP path sensor support (ACX710, ACX5448, ACX5448-M, ACX5448-D, EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-48MP, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-24T, EX4100-F-12P, EX4100-F-48T, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vMX, QFX10002-60C, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, QFX10002, QFX10008, and QFX10016)—Junos OS Release 24.2R1 supports the streaming of telemetry MAC address and MAC address and IP path data from the forwarding database to a collector using the OpenConfig resource path `/network-instances/network-instance/fdb/`. This feature is based on data models `openconfig-network-instance.yang` (version 1.2.0) and `openconfig-network-instance-l2.yang` (version 1.2.0).

[See [Junos YANG Data Model Explorer](#).]

Network Management and Monitoring

- AES-256 Encryption Algorithm Support for SNMPv3 (ACX5448, ACX5448-M, ACX5448-D, ACX710, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, QFX10002, QFX10008, and QFX10016)—Starting in Junos OS Release 24.2R1, you can configure Advanced Encryption Standard (AES) 256 algorithm for an SNMPv3 user. To configure AES-256 algorithm for an SNMPv3 user, include the `privacy-aes256` statement at the `[edit snmp v3 usm local-engine user username]` hierarchy level. AES-256 is a symmetric encryption algorithm that uses a 256-bit key to encrypt or decrypt messages and provides high-level security for protecting sensitive information.

[See [Configure SNMPv3 Encryption Type](#).]

- Clear LLDP neighbors from an interface with the gRPC Network Operations Interface (gNOI) Layer2 service (ACX710, ACX5448, ACX5448-M, ACX5448-D, EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200 and QFX5210)—Starting in Junos OS Release 24.2R1, you can execute supported Layer2 service remote procedure calls (RPCs) to perform the equivalent of the `clear lldp neighbors interface interface-name` command.

[See [gNOI Layer 2 Service](#).]

Precision Time Protocol (PTP)

Software Installation and Upgrade

- Base OS update (ACX710, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Starting in Junos OS Release 24.2R1, Junos OS uses the FreeBSD main base OS. This upgrade provides improved security and better performance. In earlier releases, Junos OS used the FreeBSD Release 12 base OS.

[See [Junos® OS Software Installation and Upgrade Guide](#).]

Additional Features

We have extended support for the following features to these platforms.

- **sFlow ingress support** (ACX710, ACX5448, ACX5448-D, and ACX5448-M)

[See [sFlow Monitoring Technology](#).]

What's Changed

IN THIS SECTION

- [EVPN | 4](#)
- [Flow-based and Packet-based Processing | 5](#)
- [General Routing | 5](#)

- Infrastructure | 7
- User Interface and Configuration | 7
- VPNs | 7

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

EVPN

- **OISM SBD bit in EVPN Type 3 route multicast flags extended community**—In EVPN Type 3 Inclusive Multicast Ethernet Tag (IMET) route advertisements for interfaces associated with the supplemental bridge domain (SBD) in an EVPN optimized intersubnet multicast (OISM) network, we now set the SBD bit in the multicast flags extended community. We set this bit for interoperability with other vendors, and to comply with the IETF draft standard for OISM, draft-ietf-bess-evpn-irb-mcast .

[See the description of the `show route table bgp.evpn.0 ?` extensive command in [CLI Commands to Verify the OISM configuration](#).]

- **Default behavior changes and new options for the easy EVPN LAG configuration (EZ-LAG) feature**—The easy EVPN LAG configuration feature now uses some new default or derived values, as follows:
 - Peer PE device `peer-id` value can only be 1 or 2.
 - You are required to configure the loopback subnet addresses for each peer PE device using the new `loopback peer1-subnet` and `loopback peer2-subnet` options at the **edit services evpn device-attribute** hierarchy level. The commit script uses these values for each peer PE device's loopback subnet instead of deriving those values on each PE device. These replace the `loopback-subnet` option at the **edit services evpn device-attribute** hierarchy level, which has been deprecated.
 - If you configure the `no-policy-and-routing-options-config` option, you must configure a policy statement called `EXPORT-LOO` that the default underlay configuration requires, or configure the new `no-underlay-config` option and include your own underlay configuration.
 - The commit script generates "notice" messages instead of "error" messages for configuration errors so you can better handle **edit services evpn** configuration issues.
 - The commit script includes the element names you configure (such as IRB instance names and server names) in description statements in the generated configuration.
 - This feature also now includes a few new options so you have more flexibility to customize the generated configuration:

- **no-underlay-config at the edit services evpn hierarchy level**—To provide your own underlay peering configuration.
- **mtu overlay-mtu and mtu underlay-mtu options at the edit services evpn global-parameters hierarchy level**—To change the default assigned MTU size for underlay or overlay packets.

[See [Easy EVPN LAG Configuration](#).]

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the **edit protocols evpn** hierarchy level. In most use cases, you don't need to change the default limit. If you want to change the default limit, we recommend that you don't set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

[See [mac-ip-limit](#).]

Flow-based and Packet-based Processing

- The subscription path for the flow sensor is changed from `/junos/security/spu/flow/usage` to `/junos/security/spu/flow/statistics`. This change maintains a uniform path in request and response data.

General Routing

- **New commit check for MAC-VRF routing instances with the encapsulate-inner-vlan statement configured**—We introduced a new commit check that prevents you from configuring an IRB interface and the `encapsulate-inner-vlan` statement together in a MAC-VRF routing instance. Please correct or remove these configurations prior to upgrading to 23.2R2 or newer to avoid a configuration validation failure during the upgrade.

[See [encapsulate-inner-vlan](#).]

- Starting in Junos OS Release 24.2R1, when you run the `run show lldp local-information interface <interface-name> | display xml` command, the output is displayed under the `lldp-local-info` root tag and in the `lldp-local-interface-info` container tag. When you run the `run show lldp local-information interface | display xml` command, the `lldp-tlv-filter` and `lldp-tlv-select` information are displayed under the `lldp-local-interface-info` container tag in the output.

- **Non-revertive switchover for sender based MoFRR**— In earlier Junos releases, source-based MoFRR ensured that the traffic reverted to the primary path from the backup path, when the primary path or session was restored. This reversion could result in traffic loss. Starting in Junos OS 22.4R3-S1, source-based MoFRR will not revert to the primary path, i.e. traffic will continue to flow through the backup path as long as the traffic flow rate on the backup path does not go below the configured threshold set under the protocols `mvpn hot-root-standby min-rate` command.
- Show active forwarding session for sender based MoFRR— The `show multicast route extensive` command will show the active forwarding session in the case of source-based MoFRR. The field Session Status: Up & Forwarding will indicate that the particular session is currently forwarding traffic.
- **Change in options and generated configuration for the EZ-LAG configuration IRB subnet-address statement**—With the EZ-LAG `subnet-address inet` or `subnet-address inet6` options at the **edit services evpn evpn-vxlan irb *irb-instance*** hierarchy, you can now specify multiple IRB subnet addresses in a single statement using the list syntax **addr1 addr2 ...** . Also, in the generated configuration for IRB interfaces, the commit script now includes default router-advertisement statements at the **edit protocols** hierarchy level for that IRB interface.

[See [subnet-address \(Easy EVPN LAG Configuration\)](#).]

- Three new VSA's have been added to code repository for 802.1x authentication on RADIUS server under Vendor ID: 2636: - 53: Event-Type - 54: Sub-Event-Type - 55: Juniper-Generic-Message

[See [Radius Attributes and VSA list supported by 802.1X](#).]

- **Change to the commit process**—In prior Junos OS and Junos OS Evolved releases, if you use the `commit prepare` command and modify the configuration before activating the configuration using the `commit activate` command, the prepared commit cache becomes invalid due to the interim configuration change. As a result, you cannot perform a regular commit operation using the `commit` command. The CLI shows an error message: 'error: Commit activation is pending, either activate or clear commit prepare'. If you now try running the `commit activate` command, the CLI shows an error message: 'error: Prepared commit cache invalid, failed to activate'. You then must clear the prepared configuration using the `clear system commit prepared` command before performing a regular commit operation. From this Junos and Junos OS Evolved release, when you modify a device configuration after 'commit prepare' and then issue a 'commit', the OS detects that the prepared cache is invalid and automatically clears the prepared cache before proceeding with regular 'commit' operation.

[See [Commit Preparation and Activation Overview](#).]

Infrastructure

- **Option to disable path MTU discovery**—Path MTU discovery is enabled by default. To disable it for IPv4 traffic, you can configure the `no-path-mtu-discovery` statement at the `edit system internet-options` hierarchy level. To reenable it, use the `path-mtu-discovery` statement.

[See [Path MTU Discovery](#).]

User Interface and Configuration

- **Configuration database maximum size increased (ACX Series, EX Series, MX Series, QFX Series, SRX Series, and vSRX)**—We've enhanced the `extend-size` statement at the `[edit system configuration-database]` hierarchy level to increase the maximum database size. On devices with a default configuration database size of ~400 MB, `extend-size` increases the maximum database size to ~2 GB. On devices with a default configuration database size of ~660 MB, `extend-size` increases the maximum database size to ~2.2 GB.

[See [configuration-database](#).]

VPNs

- **Increase in revert-delay timer range**— The `revert-delay` timer range is increased to 600 seconds from 20 seconds.
- **Configure min-rate for IPMSI traffic explicitly**— In a source-based MoFRR scenario, you can set a `min-rate` threshold for IPMSI traffic explicitly by configuring `ipmsi-min-rate` under `set routing-instances protocols mvpn hot-root-standby min-rate`. If not configured, the existing `min-rate` will be applicable to both IPMSI and SPMSI traffic.

[See [min-rate](#).]

Known Limitations

IN THIS SECTION

- [General Routing | 8](#)

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

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

General Routing

- The configuration to logout console session on disconnect does not work on ACX710. [PR1791623](#)

Open Issues

IN THIS SECTION

- [General Routing | 9](#)

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

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

General Routing

- When restart chassis-control triggered on M/MX router has configuration with ccc instance, syslog is error out " Err] ACX_ASIC_PROGRAMMING_ERROR: pfe_dnx_translation_set: Error, bcm_vlan_port_translation_set rv:Entry not found ". [PR1764966](#)
- On ACX5048 and ACX5096 platform, after the device is upgraded, disabling an interface and then rebooting the device will cause a critical issue. All interfaces will go down, resulting in a complete traffic drop. There is no known workaround to prevent this service interruption during the upgrade process. [PR1786687](#)
- On Junos OS ACX710 platform with BGP (Border Gateway Protocol) configuration, the response message will be lost and it will lead to element being stuck in the KRT (Kernel Routing Table) queue. [PR1787707](#)
- ACX1100 PTP(enterprise profile) is stuck at freerun state after upgrading junos to 21.2R3 [PR1789694](#)
- ACX710 does not recognize GPON OLT 740-124448 reports NON-JNPR after ACX power cycle. The same error state NON-JNPR can be observed when GPON OLT SFP is installed into ACX router. [PR1801112](#)
- Multicast route is reset every 5 minutes with IGMP receiver on ACX2200 with small traffic loss. Multicast route that are not active would get reset after 5 minutes due to cache timeout. This was happening even for active routes that had traffic. [PR1805017](#)

Resolved Issues

IN THIS SECTION

- [General Routing | 10](#)

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

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

General Routing

- PFE might crash on ACX710 during init. [PR1604346](#)
- We might encounter jdhcpd core during initialization. The core is rare, and there is no service impact because of this core (as the process recovers immediately). [PR1730717](#)
- ACX710 CFM asynchronous-notification feature driven on CCC-down is not supported. [PR1784447](#)
- The egress ports on ACX710 incorrectly tagging traffic expected to be untagged over CCC/VPLS interfaces. [PR1789949](#)
- PICD core can be seen during FPC is off-line/on-line, HA switch over, and system restart. [PR1793824](#)
- The l2circuit interface ccc "with Native-VLAN" configured do not add vlan-ID when receiving untag. packet [PR1793829](#)
- ACX node acting as ASBR+PE in MPLS Inter-AS Option B/C is not imposing VPN labels. [PR1794718](#)
- Port goes down after adding interface configuration and changing the port from 1g copper to 10g fiber. [PR1794939](#)
- Commit check failure will be reported when family ethernet-switching is configured on a EVPN ETREE/EVPN ELAN interface. [PR1798425](#)
- Acx-arm-feb core might be triggered if IGMP snooping is enabled and IGMP query is received on the same port as IGMP Join. [PR1799619](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 11

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

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

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

We have two types of releases, 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 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
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 cPCE

IN THIS SECTION

- [What's New | 12](#)
- [What's Changed | 13](#)
- [Known Limitations | 13](#)
- [Open Issues | 13](#)
- [Resolved Issues | 13](#)

These release notes accompany Junos OS Release 24.2R1 for cPCE. 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

- [Additional Features | 12](#)

Learn about new features introduced in this release for cPCE.

Additional Features

We have extended support for the following features to these platforms.

- **Offload constraint-based path computation onto the cPCE with cRPD support**— A containerized routing protocol process (daemon) enables containerized Path Computation Engine (cPCE) to run on-device or off-device. You can offload constraint-based path computation onto the cPCE. cPCE is

useful for computing and delegating traffic-engineered label-switched paths in an RSVP–Traffic Engineering (RSVP-TE) network deployment.

[See [cPCE Deployment Guide](#).]

What's Changed

There are no changes in behavior and syntax in this release for cPCE.

Known Limitations

There are no known limitations in hardware or software in this release for cPCE.

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

Open Issues

There are no known issues in hardware or software in this release for cPCE.

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 this release for cPCE.

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

Junos OS Release Notes for cRPD

IN THIS SECTION

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

What's New

IN THIS SECTION

- [Routing Protocols | 14](#)

Learn about new features introduced in this release for cRPD.

Routing Protocols

- **Support for BGP VPN to Global RIB Import (cRPD and MX480)**—Starting in Junos OS Release 24.2R1, we support leaking of BGP VPN routes to global RIBs to provide service providers the flexibility to allow internet access to VPN customers. To configure this feature, include the `vpn-global-import policy` statement at the `[edit routing-options inet.0]` hierarchy level.

To use the auto router discovery feature with router-id without allocating an IP-address include the `route-distinguisher-id-use-router-id` statement at the `[edit routing-options]` hierarchy level.

[See [route-distinguisher-id-use-router-id](#), and [vpn-global-import](#).]

Known Limitations

There are no known limitations in hardware or software in this release 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 or software in this release 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

There are no resolved issues in this release for cRPD.

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

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 | 17](#)

What's New

IN THIS SECTION

- [Network Management and Monitoring](#) | 16

Learn about new features introduced in this release for cSRX.

Network Management and Monitoring

- **Logging Infrastructure Support for RADIUS Accounting (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we've introduced a logging infrastructure for RADIUS accounting. The RADIUS accounting enables you to send the logs directly from dataplane to the RADIUS accounting server. When you enable RADIUS logging, logs are sent for NAT PBA ALLOC, INTERIM, and RELEASE events to the configured RADIUS server. This feature enhances the existing stream-based logging and includes:
 - Incorporation of vendor-specific attributes (VSAs) in RADIUS accounting messages
 - Support multiple RADIUS accounting servers under different streams
 - Manage retries and retransmissions of RADIUS accounting messages in case of failure
 - Flexible and capable of supporting a backup RADIUS accounting server.

To support the feature, we've introduced the following configuration statements:

- `radius`
- `retry-count`
- `radius-accounting`
- `subscriber-extension`

Use the following commands to view and to clear the RADIUS server counters for RADIUS streams:

- `show security log radius stream`
- `clear security log radius stream.`

[See [radius \(Security Log\)](#), [retry-count \(Security Log\)](#), [radius-accounting](#), and [subscriber-extension](#).]

What's Changed

IN THIS SECTION

- [Platform and Infrastructure](#) | 17

Learn about what changed in this release for cSRX.

Platform and Infrastructure

- Advanced Policy-Based Routing Policies (APBR) is not supported on cSRX instances. So, when you run the APBR related CLI commands such as `show security advance-policy-based-routing count`, then you will receive an error message as **error: Unrecognized command (network-security)**.

Known Limitations

There are no known limitations in hardware or software in this release 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 or software in this release 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 this release for cSRX.

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

Junos OS Release Notes for EX Series

IN THIS SECTION

- [What's New 24.2R1-S1 | 18](#)
- [What's New | 19](#)
- [What's Changed | 34](#)
- [Known Limitations | 37](#)
- [Open Issues | 38](#)
- [Resolved Issues | 41](#)
- [Migration, Upgrade, and Downgrade Instructions | 43](#)

What's New 24.2R1-S1

IN THIS SECTION

- [Hardware | 18](#)

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

Hardware

- **Improved PoE port power (EX4400-48MXP and EX4400-48XP)**-Starting in Junos OS Release 24.2R1-S1, the new EX4400-48MXP and EX4400-48XP switches support a PoE power budget of 3600 W. The switches support the 2000 W AC PSU and thereby higher PoE budgets of up to 3600 W with dual PSU operating in AC High voltage line. All 48 ports of the switches support PoE-bt delivering maximum value of 90 W per port.

The EX4400-48XP has the following ports:

- 10-Mbps/100-Mbps/1000-Mbps PoE ports: 48
- 100GbE QSFP28 ports: 2

The EX4400-48MXP has the following ports:

- 100-Mbps/1-Gbps/2.5-Gbps/5-Gbps/10-Gbps PoE ports: 12
- 100-Mbps/1-Gbps/2.5-Gbps PoE ports: 36
- 100GbE QSFP28 ports: 2

Both switches support the 1x100 GbE QSFP28, 4x10 GbE SFP+, and 4x25 GbE SFP28 extension modules.

[See [EX4400 Hardware Guide](#).]

What's New

IN THIS SECTION

- [Authentication and Access Control | 20](#)
- [EVPN | 20](#)
- [Interfaces | 27](#)
- [Junos OS API and Scripting | 27](#)
- [Junos Telemetry Interface | 27](#)
- [Network Management and Monitoring | 28](#)
- [Routing Policy and Firewall Filters | 29](#)
- [Routing Protocols | 29](#)
- [Serviceability | 31](#)
- [Software Installation and Upgrade | 31](#)
- [System Logging | 32](#)
- [Virtual Chassis | 32](#)
- [Additional Features | 32](#)

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

Authentication and Access Control

- **RADSEC support**—Starting in Junos OS release 24.2R1, RADSEC is supported for ACX, EX, MX, and SRX switches. The RADSEC protocol provides secure transport of RADIUS authentication and accounting data across untrusted networks using Transport Layer Security (TLS) over TCP as the transport protocol.

EVPN

- **Default discard policy for GBP filters (EX4100, EX4400, EX4650, and QFX5120)**—Starting in Junos OS Release 24.2R1, you can configure group-based policy (GBP) firewall filters with a default discard policy that is applicable when a packet fails to meet any of the match conditions.

[See [Example: Micro and Macro Segmentation Using Group Based Policy in a VXLAN](#).]

- **MAC/IP inter-tagging for GBP filters (EX4100, EX4400, EX4650, and QFX5120)**—Starting in Junos OS Release 24.2R1, you can apply media access control (MAC)-based GBP firewall filters to routed traffic and IP-based GBP firewall filters to switched traffic. This is called inter-tagging. Previously, MAC-based GBP filters applied to switched traffic and IP-based GBP filters applied to routed traffic. By enabling inter-tagging, your MAC-based and IP-based GBP filters apply to both switched and routed traffic.

[See [Example: Micro and Macro Segmentation Using Group Based Policy in a VXLAN](#).]

- **Ingress policy enforcement and tag propagation (EX9204, EX9208, and EX9214)**—Starting in Junos OS Release 24.2R1, the EX9204, EX9208, and EX9214 switches support ingress policy enforcement of group-based policy (GBP) firewall filters and GBP tag propagation for /32 IP routes. Ingress policy enforcement and GBP tag propagation save network bandwidth by discarding tagged packets at the ingress that would otherwise be discarded at the egress.

[See [Example: Micro and Macro Segmentation Using Group Based Policy in a VXLAN](#).]

- **GBP tag propagation using EVPN Type 5 route advertisements (EX4400, EX4650, and QFX5120)**—Starting in Junos OS Release 24.2R1, we support group-based policy (GBP) tag propagation using EVPN Type 5 route advertisements of IP prefixes. Switches and routers typically use EVPN Type 5 advertisements for exchanging routes between data centers. Prior to this release, we supported EVPN Type 2 to Type 5 route conversion between data centers, which resulted in /32 IP routes being exchanged instead of IP prefix routes.

[See [Example: Micro and Macro Segmentation Using Group Based Policy in a VXLAN](#).]

- **Access security support in EVPN-VXLAN overlay networks (EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, we support access security features on certain EX Series and QFX Series switches that function as Layer 2 VXLAN gateways in an Ethernet VPN–Virtual Extensible LAN (EVPN-

VXLAN) centrally-routed overlay network (two-layer IP fabric). We support the following features on Layer 2 server-facing interfaces that are associated with VXLAN-mapped VLANs:

- DHCPv4 and DHCPv6 snooping [See [DHCP Snooping](#).]
- Dynamic ARP inspection (DAI) [See [Understanding and Using Dynamic ARP Inspection \(DAI\)](#).]
- Neighbor discovery inspection (NDI) [See [IPv6 Neighbor Discovery Inspection](#).]
- IPv4 and IPv6 source guard [See [Understanding IP Source Guard for Port Security on Switches](#).]
- Router advertisement (RA) guard [See [Understanding IPv6 Router Advertisement Guard](#).]

The access security features function the same and you configure them in the same way in an EVPN-VXLAN environment as you do in a non-EVPN-VXLAN environment. However, keep these differences in mind:

- We do not support these features on multihomed servers.
- These features do not influence the VXLAN tunneling and encapsulation process.
- **MAC-VRF instances with EVPN-VXLAN (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, and EX4300-MP)**—Support for the `mac-vrf` instance type includes `vlan-based`, `vlan-aware`, and `vlan-bundle` service types for EVPN unicast routes only.

[See [MAC-VRF Routing Instance Type Overview](#), [mac-vrf](#), and [service-type](#).]

- **MAC-VRF with EVPN-VXLAN (EX9204 and EX9208 switches)**—Data center service providers must support multiple customers with their own routing and bridging policies in the same physical network. To accommodate this requirement, you can now configure multiple customer-specific EVPN instances (EVIs) of type `mac-vrf`, each of which can support a different EVPN service type. This configuration results in customer-specific virtual routing and forwarding (VRF) tables with MAC addresses on each Juniper Networks device that serves as a virtual tunnel endpoint (VTEP) in the Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) network.



NOTE: We support MAC-VRF routing instances for EVPN unicast routes only.

To support this feature, we introduce a uniform routing instance configuration that complies with RFC 7432, **BGP MPLS-Based Ethernet VPN**. The uniform configuration eliminates hardware restrictions that limit the number of EVIs and the combinations of EVIs with their respective policies that can simultaneously exist. The common configuration includes the following new CLI elements:

- The `mac-vrf` keyword at the `[edit routing-instances name instance-type]` hierarchy level.

- The service-type configuration statement at the `[edit routing-instances name]` hierarchy level. We support VLAN-based, VLAN-aware, and VLAN-bundle service types.
- (QFX10000 line of switches only) The forwarding-instance configuration statement at the `[edit routing-instances name]` hierarchy level. With this optional configuration statement, you can map multiple routing instances to a single forwarding instance. If you don't include this configuration statement, the default forwarding instance is used.

We continue to support the existing method of routing instance configuration along with the new uniform routing instance configuration.

[See [EVPN User Guide](#).]

- **sFlow support for EVPN-VxLAN multicast (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-12P, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, and EX4400-48T)**—Starting in Junos OS Release 24.2R1, you can use the sFlow technology to sample EVPN-VxLAN multicast traffic configured at the interface level.

We support sampling with collector to be reachable through normal Layer 3 (L3) gateway (underlay), management IP (reachable through default or non-default routing-instance), or VXLAN tunnel (overlay).

To enable known multicast sampling, use the CLI command `set forwarding-options sFlow egress-multicast enable`.

[See [Overview of sFlow Technology](#).]

- **EVPN-VXLAN fabric with an IPv6 underlay (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-24P, and EX4100-24T)**—Starting in Junos OS Release 24.2R1, you can configure an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) fabric with an IPv6 underlay. You can use this feature only with MAC-VRF routing instances (all service types). You must configure either an IPv4 or an IPv6 underlay across the EVPN instances in the fabric; you can't mix IPv4 and IPv6 underlays in the same fabric.

To enable this feature, perform these steps when you configure the EVPN underlay:

- Configure the underlay VXLAN tunnel endpoint (VTEP) source interface as an IPv6 address:

```
set routing-instances mac-vrf-instance-name vtep-source-interface lo0.0 inet6
```

- Even though the underlay uses the IPv6 address family, for BGP handshaking to work in the underlay, you must configure the router ID in the routing instance with an IPv4 address:

```
set routing-instances mac-vrf-instance-name routing-options router-id ipv4-address
```

We support the following EVPN-VXLAN features with an IPv6 underlay:

- EVPN Type 1, Type 2, Type 3, Type 4, and Type 5 routes [See [EVPN Type-5 Route with VXLAN Encapsulation for EVPN-VXLAN](#).]
- IPv6 Underlay Overview [See [EVPN-VXLAN with an IPv6 Underlay](#).]
- Shared VTEP tunnels (required with MAC-VRF instances)
- All-active multihoming [See [EVPN Multihoming Overview](#).]
- EVPN core isolation [See [Understanding When to Disable EVPN-VXLAN Core Isolation](#).]
- Bridged overlays [See [Bridged Overlay Design and Implementation](#).]
- Layer 3 gateway functions in edge-routed bridging (ERB) and centrally routed bridging (CRB) overlays with IPv4 or IPv6 traffic
- Underlay and overlay load balancing
- Layer 3 protocols over integrated routing and bridging (IRB) interfaces—BFD, BGP, OSPF [See [Supported Protocols on an IRB Interface in EVPN-VXLAN](#).]
- EVPN proxy Address Resolution Protocol (ARP) and ARP suppression, and proxy Neighbor Discovery Protocol (NDP) and NDP suppression [See [EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression](#).]

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#) and [EVPN User Guide](#).]

- **L2PT with Q-in-Q over VXLAN tunnels in EVPN-VXLAN bridged overlay networks (EX4100-48P, EX4400-48F, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, we support Layer 2 protocol tunneling (L2PT) with Q-in-Q for traffic from an access interface to VXLAN tunnel destinations in a bridged overlay (BO) Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) network. You can use this feature with service-provider style or enterprise-style access interface configurations.

You can use L2PT over VXLAN tunnels with all of the Q-in-Q use cases we support for an EVPN-VXLAN network. For Q-in-Q, the device tunnels tagged frames over VXLAN using the VNI of the VLAN in the frame, and tunnels untagged frames using the VNI of the native VLAN.

To enable this feature, configure the `l2pt` statement at the `[edit protocols layer2-control]` hierarchy level with the access interface name `interface name` and the following required options:

- `destination vxlan-tunnel`—Enable L2PT for traffic toward a VXLAN tunnel destination.
- `protocol protocol-name`—Specify a protocol to tunnel. Include additional `protocol` statements for each protocol you want to tunnel.

[See [Layer 2 Protocol Tunneling over VXLAN Tunnels in EVPN-VXLAN Bridged Overlay Networks](#), [Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network](#), and [I2pt \(Destination Tunnels\)](#).]

- **Suppress EVPN Type 5 host routes from DCI to DC (EX4400-24MP, EX4400-48F, MX304, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can suppress EVPN Type 5 host route advertisements that re-originate from the data center interconnect (DCI) to the local DC. You can achieve better scaling and performance on leaf devices with this feature.

[See [suppress-host-routes-from-dci-to-dc](#).]

- **Enhanced OISM in EVPN-VXLAN ERB overlay networks with an IPv6 underlay (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can configure enhanced optimized intersubnet multicast (OISM) for IPv4 and IPv6 multicast data traffic with an Ethernet VPN–Virtual Extensible LAN (EVPN–VXLAN) edge-routed bridging (ERB) overlay network that has an IPv6 underlay. To configure this feature:
 - Set up the EVPN-VXLAN fabric with an IPv6 underlay:
 - You can use either external BGP (EBGP) or OSPFv3 with IPv6 addressing for the IPv6 underlay.
 - Use the `inet6` option when you set the VXLAN tunnel endpoint (VTEP) source interface to the device loopback interface in the EVPN instance (EVI):

```
set routing-instances evpn-instance-name vtep-source-interface lo0.0 inet6
```

- Configure the enhanced OISM elements for your multicast EVPN-VXLAN environment in the same way you would configure these elements in an EVPN-VXLAN network with an IPv4 underlay.

You can configure any of the supported platforms as enhanced OISM server leaf devices, and only EX4650 and QFX5120 switches as enhanced OISM border leaf devices.

[See [EVPN-VXLAN with an IPv6 Underlay](#) and [Optimized Intersubnet Multicast in EVPN Networks](#).]

- **Statically identify multihoming peer OISM leaf devices in an EVPN-VXLAN network running enhanced OISM (EX4100-24T, EX4300-MP, EX4400-24MP, EX4400-24P, EX4400-48F, EX4400-48MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can statically configure a multihoming peer leaf device to identify its peers in an Ethernet VPN–Virtual Extensible

LAN (EVPN-VXLAN) network running enhanced optimized intersubnet multicast (OISM). EVPN-VXLAN provider edge (PE) leaf devices are multihoming peers when they share an Ethernet segment (ES) for a multihomed host. With enhanced OISM, if the leaf devices have static information about their multihoming peers, they can avoid multicast traffic loss when their peer devices go down and up again.

On each multihoming peer leaf device, to identify the device's multihoming peers, configure the `multihoming-peer-gateways [peer-device-IPv4-address ...]` statement at the `[edit protocols evpn]` hierarchy level. Specify a list of peer addresses within square brackets, or specify a single peer address without any brackets.

[See [Statically Identify Multihoming Peers With Enhanced OISM To Improve Convergence.](#)]

- **Non-revertive preference-based DF election in EVPN-MPLS networks (EX4400-24P, MX960, and vMX)**—Starting in Junos OS Release 24.2R1, you can configure non-revertive preference-based designated forwarder (DF) election for an Ethernet segment identifier (ESI) in an Ethernet VPN—MPLS (EVPN-MPLS) network. By default, preference-based DF election for an Ethernet segment identifier (ESI) is revertive, which means:
 - If the EVPN provider edge (PE) device currently in the DF role goes down, the next preferred PE device becomes the new DF.
 - When the old DF comes back up, the DF role reverts to the old DF.

Changing the current DF role for an ESI frequently can impact traffic flow. To avoid revertive DF role changes, you can now set the non-revertive option at the `[edit interfaces name esi df-election-type preference]` hierarchy level.

We also provide new options you can configure to load-balance DF election per EVPN instance (EVI) or per ESI based on the lowest configured preference value or the highest configured preference value, as follows:

- At the EVI level—Use the `designated-forwarder-preference-least` option or the `designated-forwarder-preference-highest` option at the `[edit routing-instances evpn-instance-name protocols evpn]` hierarchy level.
- At the ESI level—Use the `least` option at the `[edit interfaces interface-name esi df-election-type preference]` or `[edit protocols evpn interconnect esi df-election-type preference]` hierarchy level.

[See [df-election-type](#) and [evpn](#)].

- **EVPN-VXLAN DCI multicast support with enhanced OISM (EX4400-24MP, EX4400-24P, EX4400-48F, EX4400-48MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, we support Ethernet VPN—Virtual Extensible LAN (EVPN-VXLAN) to EVPN-VXLAN seamless Data Center Interconnect (DCI) with enhanced optimized intersubnet multicast (OISM). Without this feature, the DCI gateway devices flood multicast traffic across the interconnecting WAN. Flooding consumes significant WAN

bandwidth if your network has many multicast flows or high multicast traffic rates. This feature seamlessly replaces the multicast flooding behavior. To optimize multicast forwarding across a DCI, OISM leaf devices in each network:

- Advertise selective multicast Ethernet tag (SMET) routes (EVPN Type 6 routes) when a receiver subscribes to a multicast flow. The DCI gateways seamlessly propagate the SMET routes across the DCI on the OISM SBD.
- Send multicast traffic based on the received SMET routes only to the remote receivers across the DCI who subscribed to that multicast flow.

To configure this feature:

- Configure the DCI gateway devices the same way you would configure the devices without multicast support.
- Configure enhanced OISM in the networks on both sides of the DCI.
 - With enhanced OISM, you can configure each OISM device with only the VLANs that the device hosts, except on multihoming peer OISM devices and the peer DCI gateways in each data center network. On those peer devices, you must configure the OISM revenue VLANs symmetrically.
 - Configure the same OISM supplemental bridge domain (SBD) in the matching tenant virtual routing and forwarding (VRF) instances on both sides of the DCI.

As OISM devices, the DCI gateways follow the enhanced OISM operational model to forward traffic to other OISM devices in their own network or across the DCI. They send the multicast traffic:

- Across the DCI to the other DCI gateways only on the OISM SBD.
- To other non-multihoming peer provider edge (PE) devices in their network only on the OISM SBD.
- To their multihoming peer PE devices only on the source VLAN.

You can also configure a DCI gateway as an OISM PIM EVPN gateway (PEG). The device acts as a DCI gateway and also as an OISM PEG border leaf device to exchange multicast traffic with devices outside of either network.

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

- **Generation of EVPN Type 3 routes on 802.1X dynamically mapped interfaces (EX4100-24MP, EX4300-MP, EX4400-24P, QFX5120-32C, QFX5120-48T, and QFX5120-48Y)**—Starting in Junos OS Release 24.2R1, we support generating Ethernet VPN (EVPN) Type 3 routes across interfaces dynamically mapped by the 802.1X protocol to a Virtual Extensible LAN (VXLAN) extended bridge domain (BD).

[See [dot1x](#).]

Interfaces

- **Support for PoE port bounce (EX Series)**—Starting in Junos OS Release 24.2R1, the EX4100 and EX4400 series switches support PoE port bounce. We have enhanced the `request interface bounce` command by adding the `poe` parameter. Run the command to reboot the PoE port with a PoE port bounce. This action eliminates the need to run the `commit` command twice to reboot the PoE port. You can use the `interval` optional parameter to set the time interval, which can be between 0 to 60 seconds, within which the port reboots. By default, the port reboots immediately.

[See [Request interface bounce](#).]

Junos OS API and Scripting

- **Support for configuring the `allow-transients` statement for individual commit scripts (EX4100-24MP, EX4400-24MP, and QFX5120-32C)**—Starting in Junos OS Release 24.2R1, you can configure the `allow-transients` statement for individual commit scripts. Configuring the `allow-transients` statement for individual scripts enables you to add transient configuration changes to the checkout configuration for specific commit scripts while still keeping transient changes disabled for the other commit scripts.

[See [allow-transients](#).]

Junos Telemetry Interface

- **OpenConfig MAC address and MAC address and IP path sensor support (ACX710, ACX5448, ACX5448-M, ACX5448-D, EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-48MP, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-24T, EX4100-F-12P, EX4100-F-48T, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vMX, QFX10002-60C, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, QFX10002, QFX10008, and QFX10016)**—Junos OS Release 24.2R1 supports the streaming of telemetry MAC address and MAC address and IP path data from the forwarding database to a collector using the OpenConfig resource path `/network-instances/network-instance/fdb/`. This feature is based on data models `openconfig-network-instance.yang` (version 1.2.0) and `openconfig-network-instance-l2.yang` (version 1.2.0).

[See [Junos YANG Data Model Explorer](#).]

- **Hardware resource threshold monitoring for capacity planning (EX4100-48MP, EX4400-24MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM0)**—Junos OS Release 24.2R1 supports hardware resource threshold monitoring. Using this feature, you can monitor

hardware resource utilization. Advance knowledge about resource utilization nearing or crossing a maximum capacity provides time for you to act and prevent network disruption and traffic loss.

Use the system packet-forwarding-options hw-resource-monitor resource-list configuration statement at the [edit] hierarchy level to create a list of hardware resources that you want to monitor. Once configured, periodic resource monitoring occurs at the polling interval you set.

View the monitored data using operational mode commands or use Junos Telemetry interface (JTI) to send data from your device to a collector using the resource path `/junos/system/linecard/npu/memory/`.

[See [Configure Hardware Threshold Monitoring for Capacity Planning](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **Chassis and event telemetry sensor support (EX4100-24MP, EX4100-24P, EX4100-48MP, EX4100-48P, EX4100-24T, EX4100-48T, EX4100-F-24P, EX4100-F-48P, EX4100-F-12P, EX4100-F-12T, EX4100-F-24T, EX4100-F-48T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, and EX4400-48T)**—Junos OS Release 24.2R1 supports the Routing Engine chassis-related sensors under the resource path `/components/component/` and the events data sensors under the native resource path `/junos/events/event/`. You can stream data from a device to a collector using native (UDP) sensors, Juniper proprietary Remote Procedure Call (gRPC) service, or gRPC Network Management Interface (gNMI).
[For sensors, see [Junos YANG Data Model Explorer](#).]
- **Native support for interfaces and chassis sensors (EX9204, EX9208, and EX9214)**—Junos OS Release 24.2R1 supports native streaming of operational state statistics and counters for chassis and interfaces using new Junos-specific sensor paths. To stream data, use the sensors `/state/chassis/` and `/state/interfaces/`.

[See [Junos YANG Data Model Explorer](#).]

Network Management and Monitoring

- **Clear LLDP neighbors from an interface with the gRPC Network Operations Interface (gNOI) Layer2 service (ACX710, ACX5448, ACX5448-M, ACX5448-D, EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200 and QFX5210)**—Starting in Junos OS Release 24.2R1, you can execute supported Layer2 service remote procedure calls (RPCs) to perform the equivalent of the `clear lldp neighbors interface interface-name` command.

[See [gNOI Layer 2 Service](#).]

Routing Policy and Firewall Filters

- **Support to configure DDoS protocol using CLI (EX4100 and EX4400)**—Starting in Junos OS Release 24.2R1, you can configure the distributed denial of service (DDoS) protocol using CLI on EX4100 and EX4400 devices. You can also use the following operational commands to view the DDoS protocol information:

- `show ddos-protection protocols`
- `show ddos-protection statistics`
- `show ddos-protection protocols violations`
- `show ddos-protection protocols parameters`
- `show ddos-protection protocols statistics`
- `clear ddos-protection protocols`

[See [ddos-protection \(DDoS\)](#), [show ddos-protection protocols](#), [clear ddos-protection protocols](#), [show ddos-protection statistics](#), [show ddos-protection protocols violations](#), [show ddos-protection protocols parameters](#), and [show ddos-protection protocols statistics](#).]

Routing Protocols

- **Support for OSPFv2 HMAC SHA-2 keychain authentication and weighted ECMP (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-48MP, EX4100-H-12P, EX4100-H-12P-DC, EX4100-H-24P, EX4100-H-24P-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX304, MX150, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, and VMX)**—Starting in Junos OS Release 24.2R1, you can enable OSPFv2 keychain module with HMAC-SHA2 authentication to authenticate packets reaching or originating from an OSPF interface. HMAC SHA2 algorithms include HMAC-SHA2-256, HMAC-SHA2-384 and HMAC-SHA2-512 as defined in RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*. We also support the HMAC-SHA2-224 algorithm. This feature ensures smooth transition from one key to another for OSPFv2 with enhanced security. We also support HMAC-SHA1 and HMAC-SHA2 authentication for virtual and sham links.

You can enable weighted ECMP for directly connected routers. In earlier releases, Junos OS ECMP algorithm does not take the underlying bandwidth into consideration. The algorithm assumes that the links are of equal capacity and the traffic is distributed equally based on this assumption.

To enable OSPFv2 HMAC-SHA2 authentication, configure the keychain *keychain-name* configuration statement [edit protocols ospf area *area-id* interface *interface-name* authentication] at the hierarchy level and algorithm (hmac-sha2-224 | hmac-sha2-256 | hmac-sha2-384 | hmac-sha2-512) option at the [edit security authentication-key-chains key-chain *key-chain-name*] hierarchy level.

To enable keychains authentication support for OSPFv2 virtual links, configure the keychain *keychain-name* configuration statement [edit protocols ospf area *area-id* virtual-link *neighbor-id* router-id transit-area *area-id* authentication] at the hierarchy level.

To enable keychains authentication support for OSPFv2 sham links, configure the keychain *keychain-name* configuration statement [edit protocols ospf area *area-id* virtual-link *neighbor-id* router-id transit-area *area-id* authentication] at the hierarchy level.

To enable weighted ECMP traffic distribution on directly connected OSPFv2 neighbors, configure weighted one-hop statement at the [edit protocols ospf spf-options multipath] hierarchy level.

[See [Understanding OSPFv2 Authentication](#) and [Understanding Weighted ECMP Traffic Distribution on One-Hop OSPF Neighbors](#) .]

- **BGP link bandwidth community (cRPD, EX4100-48MP, EX4300-MP, EX4400-48MP, EX4650, EX9204, EX9208, MX240, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020, cSRX, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200, and QFX5210)**—Starting in Junos OS Release 24.2R1, BGP can communicate link speeds to remote peers, enabling better optimization of traffic distribution for load balancing. A BGP group can send the *link-bandwidth* non-transitive extended community over an EBGp session for originated or received and readvertised link-bandwidth extended communities.

To configure the non-transitive link bandwidth extended community, include the bandwidth-non-transitive: *value* in the export policy at the [edit policy-options community *name* members *community-ids*] hierarchy level.

To enable the device to automatically detect and attach the link-bandwidth community on a route at import, include the [auto-sense](#) auto-sense statement at the [edit protocols bgp group link-bandwidth] hierarchy level. This feature facilitates the integration of devices with different transmission speeds within the network, enabling efficient traffic distribution based on link speed.

[See and [group \(Protocols BGP\)](#).]

- **HMAC authentication with hash functions for IS-IS (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-48MP, EX4100-H-12P, EX4100-H-12P-DC, EX4100-H-24P, EX4100-H-24P-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4600-VC, EX4650, EX4650-48Y VC, EX9204, EX9208, EX9214, MX204, MX240, MX304, MX150, MX480, MX960,**

MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)—Starting in Junos OS Release 24.2R1, we extend support to the IS-IS keychain with the following hash functions:

- HMAC-SHA2-224,
- HMAC-SHA2-256,
- HMAC-SHA2-384,
- HMAC-SHA2-512

Currently, IS-IS supports inline authentication using simple password, keyed MD5 and HMAC-SHA1 algorithms with common keychain. Note that it's important to have the system time synchronized on all nodes when a keychain is active on an IS-IS session.

[See [Understanding Hitless Authentication Key Rollover for IS-IS.](#)]

Serviceability

- **Support for automatic Junos volume recovery using OAM (EX Series)**—Starting in Junos OS Release 24.2R1, we've introduced automatic recovery functionality for the Junos volume. To enable this feature, configure the `automatic recovery` option under the `[edit system snapshot]` hierarchy level.

By enabling the `automatic recovery` option under `[edit system snapshot]` hierarchy and having a recovery snapshot available, the device automatically recovers the Junos volume when the device can not load the Junos volume. Junos volume recovery happens automatically from the OAM volume and then boots to the Junos volume.

This feature simplifies the recovery process and ensures that the device boots to the Junos volume without any manual intervention.

[See [recovery](#) and [snapshot](#).]

Software Installation and Upgrade

- **Base OS update (ACX710, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, Junos OS uses the FreeBSD main base OS. This upgrade provides improved security and better performance. In earlier releases, Junos OS used the FreeBSD Release 12 base OS.

[See [Junos® OS Software Installation and Upgrade Guide](#).]

- **In-band ZTP management in campus fabrics (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX9204, EX9208, EX9214, MX304, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can simplify the provisioning process for campus fabrics. Already provisioned

upstream devices, such as core and distribution devices, that are capable of detecting downstream Day 0 devices can provide Layer 3 connectivity. With Layer 3 connectivity, the downstream Day 0 devices can proceed with Secure ZTP.

To configure in-band ZTP management, enable the `in-band-ztp` statement at the `[edit system services]` hierarchy on your core and distribution devices. Optionally, your cloud controller can provide the in-band-ztp configuration as part of the provisioning process for your core and distribution devices.

See [Zero Touch Provisioning](#)

System Logging

- **Support for sending system log messages from the default routing instance when the dedicated management instance is configured (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-48MP, EX4100-H-12P, EX4100-H-12P-DC, EX4100-H-24P, EX4100-H-24P-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208 and EX9214)**—Starting in Junos OS Release 24.2R1, when you configure the system logging information, the system does not have to use the dedicated management instance. System logging traffic prioritizes routing instances configured at the host level. Next, it prioritizes the routing instances configured at the `syslog` level. If you do not configure a routing instance at either of these levels, even if the management instance is configured at the global level, the system log messages default to the default routing instance and the `inet.0` routing table. Prior to this release, when the dedicated management instance `mgmt_junos` was configured, system logging traffic used it by default. Thus, system logs reach the host only if the host is reachable by the default `inet.0` routing instance.

[See [System Logging and Routing Instances](#) and [Management Interface in a Dedicated Instance](#).]

Virtual Chassis

- **Enhanced jfirmware-based firmware upgrade support on Virtual Chassis (EX4400, EX4100)**—Starting in Junos OS Release 24.2R1, enhanced jfirmware based firmware upgrade support is offered on Virtual Chassis for EX4400 and EX4100 switches. The components covered as part of proposed enhancements are BIOS, CPU CPLD, and SYS CPLD.
- **Virtual Chassis (HGoE) Support** —Starting in Junos OS Release 24.2R1, support is added for Virtual Chassis using HiGig over Ethernet (HGoE Mode) on EX4100/EX4100-F switches.

[See [Understanding EX Series Virtual Chassis](#).]

Additional Features

We have extended support for the following features to these platforms.

- **RPM and TWAMP support** (EX4100, EX4400, and EX4650 switches). Starting in Junos OS Release 24.2R1, these EX Series switches fully support RPM and Two-Way Active Measurement Protocol (TWAMP):
 - You can configure RPM probe generation, reception, and reflection and enable timestamps on RPM probe messages.
 - You can also use RPM probes to detect link status and to change the preferred-route state on the basis of the probe results. RPM-tracked routes can be IPv4 or IPv6, and support up to 16 next hops for each IPv4 or IPv6 RPM-tracked static route. You can configure route preference and tag values for each IPv4 or IPv6 destination prefix.
 - You can also configure RPM services to automatically determine whether a path exists between a host device and its configured BGP neighbors. You can view the results of the discovery using an SNMP client. The results are stored in `pingResultsTable`, `jnxPingResultsTable`, `jnxPingProbeHistoryTable`, and `pingProbeHistoryTable`.
 - For TWAMP, we support IPv4 and IPv6 (including link-local addresses) traffic for control sessions and test sessions.
 - For TWAMP Light, we support IPv4 and IPv6 traffic, including IPv6 link-local addressing.

[See [Understanding Real-Time Performance Monitoring on EX and QFX Switches](#), [rpm-tracking](#), [show route rpm-tracking](#), and [Understand Two-Way Active Measurement Protocol](#).]

- **Seamless EVPN-VXLAN stitching** (EX4100-48MP, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T). We support the seamless stitching of unicast and broadcast, unknown unicast, and multicast (BUM) routes in an interconnected Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) environment.

[See [interconnect](#).]

- **Supported transceivers, optical interfaces, and DAC cables** Select your product in the Hardware Compatibility Tool (<https://apps.juniper.net/hct/product/>) to view supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.
- **Supported transceivers, optical interfaces, and DAC cables** Select your product in the Hardware Compatibility Tool (<https://apps.juniper.net/hct/product/>) to view supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.
- **Support for enhanced firewall filter processing using hardware-assisted segmented filters for large filters** (EX9200)

[See [fast-lookup-filter](#).]

- **Support for TCAM group optimization** (EX4100-48MP, EX4100-H-12P, EX4100-H-12P-DC, EX4100-H-24F, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, and EX4400-48T)

[See [loopback-firewall-optimization](#).]

- **Wake-on LAN targeted-broadcast feature for EVPN-VXLAN networks** (EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4100-F-12P, EX4100-F-12T, EX4100-F-24P, EX4100-F-24T, EX4100-F-48P, EX4100-F-48T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)

[See [Targeted Broadcast](#) and [targeted-broadcast](#).]

What's Changed

IN THIS SECTION

- [General Routing | 34](#)
- [EVPN | 35](#)
- [User Interface and Configuration | 37](#)
- [VPN | 37](#)

Learn about what changed in this release for EX Series switches.

General Routing

- **Change in options and generated configuration for the EZ-LAG configuration IRB subnet-address statement**--With the EZ-LAG subnet-address inet or subnet-address inet6 options at the [edit services evpn evpn-vxlan irb <varname>irb-instance</varname>] hierarchy, you can now specify multiple IRB subnet addresses in a single statement using the list syntax [addr1 addr2 ...]. Also, in the generated configuration for IRB interfaces, the commit script now includes default router-advertisement statements at the [edit protocols] hierarchy level for that IRB interface.

[See [subnet-address \(Easy EVPN LAG Configuration\)](#).]

- **Change in use of RSA signatures with SHA-1 hash algorithm**—Starting in Junos OS Release 24.2R1, there is a behavioural change by OpenSSH 8.8/8.8p1. OpenSSH 8.8/8.8p1 disables the use of RSA signatures with SHA-1 hash algorithm by default. You can use RSA signatures with SHA-256 or SHA-512 hash algorithm.

[See [show multicast route](#).]

- **Show active forwarding session for sender based MoFRR**-- The `show multicast route extensive` command will show the active forwarding session in the case of source-based MoFRR. The field `Session Status: Up & Forwarding` will indicate that the particular session is currently forwarding traffic.
- When you run the `run show lldp local-information interface <interface-name> | display xml` command, the output is displayed under the `lldp-local-info` root tag and in the `lldp-local-interface-info` container tag. When you run the `run show lldp local-information interface | display xml` command, the `lldp-tlv-filter` and `lldp-tlv-select` information are displayed under the `lldp-local-interface-info` container tag in the output.
- Starting in Junos OS Release 24.2R1, when you run the `run show lldp local-information interface interface-name | display xml` command, the output is displayed under the `lldp-local-info` root tag and in the `lldp-local-interface-info` container tag. When you run the `run show lldp local-information interface | display xml` command, the `lldp-tlv-filter` and `lldp-tlv-select` information are displayed under the `lldp-local-interface-info` container tag in the output.
- **Show active forwarding session for sender based MoFRR**— The `show multicast route extensive` command will show the active forwarding session in the case of source-based MoFRR. The field `Session Status: Up and Forwarding` will indicate that the particular session is currently forwarding traffic.

See [show multicast route](#)

EVPN

- **OISM SBD bit in EVPN Type 3 route multicast flags extended community**—In EVPN Type 3 Inclusive Multicast Ethernet Tag (IMET) route advertisements for interfaces associated with the supplemental bridge domain (SBD) in an EVPN optimized intersubnet multicast (OISM) network, we now set the SBD bit in the multicast flags extended community. We set this bit for interoperability with other vendors, and to comply with the IETF draft standard for OISM, `draft-ietf-bess-evpn-irb-mcast`.

See the description of the `show route table bgp.evpn.0 ? extensive` command in [CLI Commands to Verify the OISM configuration](#).

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the `[edit protocols evpn]` hierarchy level. In most use cases, you don't need to change the default limit. If you want to change the default limit, we recommend that you don't set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

[See [mac-ip-limit](#).]

- **Default behavior changes and new options for the easy EVPN LAG configuration (EZ-LAG) feature**—The easy EVPN LAG configuration feature now uses some new default or derived values, as follows:
 - Peer PE device `peer-id` value can only be 1 or 2.
 - You are required to configure the loopback subnet addresses for each peer PE device using the new `loopback peer1-subnet` and `loopback peer2-subnet` options at the **edit services evpn device-attribute** hierarchy level. The commit script uses these values for each peer PE device's loopback subnet instead of deriving those values on each PE device. These replace the `loopback-subnet` option at the **edit services evpn device-attribute** hierarchy level, which has been deprecated.
 - If you configure the `no-policy-and-routing-options-config` option, you must configure a policy statement called `EXPORT-LO0` that the default underlay configuration requires, or configure the new `no-underlay-config` option and include your own underlay configuration.
 - The commit script generates "notice" messages instead of "error" messages for configuration errors so you can better handle **edit services evpn** configuration issues.
 - The commit script includes the element names you configure (such as IRB instance names and server names) in description statements in the generated configuration.

This feature also now includes a few new options so you have more flexibility to customize the generated configuration:

- `no-underlay-config` at the **edit services evpn** hierarchy level—To provide your own underlay peering configuration.
- `mtu overlay-mtu` and `mtu underlay-mtu` options at the **edit services evpn global-parameters** hierarchy level—To change the default assigned MTU size for underlay or overlay packets.

See [Easy EVPN LAG Configuration](#).

- **Group-based Policy (GBP) tag displayed with CLI command**—On platforms that support VXLAN-GBP, the `show bridge mac-table` command now displays a GBP TAG output column that lists the GBP tag associated with the MAC address for a bridge domain or VLAN in a routing instance. Even if the device doesn't support or isn't using GBP itself, the output includes this information for GBP tags in packets received from remote EVPN-VXLAN peers.

See [Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN](#)

User Interface and Configuration

- **Configuration database maximum size increased (ACX Series, EX Series, MX Series, QFX Series, SRX Series, and vSRX)**—We've enhanced the extend-size statement at the [edit system configuration-database] hierarchy level to increase the maximum database size. On devices with a default configuration database size of ~400 MB, extend-size increases the maximum database size to ~2 GB. On devices with a default configuration database size of ~660 MB, extend-size increases the maximum database size to ~2.2 GB.

[See [configuration-database](#).]

VPN

- **Increase in revert-delay timer range**--The revert-delay timer range is increased to 600 seconds from 20 seconds.

[See [min-rate](#).]

Known Limitations

IN THIS SECTION

- [General Routing](#) | 38

Learn about known limitations in this release for EX Series switches.

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

General Routing

- On EX2300, EX3400, EX4300-48MP and EX4300 , Pause frames counters does not get incremented when pause frames are sent. [PR1580560](#)

Open Issues

IN THIS SECTION

- [General Routing | 38](#)
- [Infrastructure | 40](#)
- [Interface and Chassis | 40](#)
- [Platform and Infrastructure | 40](#)

Learn about open issues in this release for EX Series switches.

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

General Routing

- Runt, fragment and jabber counters are not incrementing on EX4300-MPs. [PR1492605](#)
- On all EX platforms, whenever beacon LED functionality is enabled, there is a mismatch between the physical LED status and the output of the CLI command `show chassis led` showing incorrect port LED status for interfaces as LED up instead of off. [PR1697678](#)
- On EX2300-48 MP without VC pre-provisioned configuration, If primary's member-id and primary member's interface configuration are changed then VC is taking more time to get stabilized. [PR1764542](#)
- During device reboot, MGE connected ports on the peer goes up after 90s into reboot. [PR1767347](#)

- After rebooting a mixed Virtual Chassis (VC) of EX4300-P and EX4300-MP switches or rebooting a EX4300-P member, interfaces with Power over Ethernet (PoE) configured won't come up on EX4300-P members. [PR1782445](#)
- Ex-Hardening:Local/Remote fault insertion from TG is failing. [PR1789999](#)
- (EX4400): after clearing arp on access error seen on aggregation devices with metro configuration.[PR1793885](#)
- On EX platforms, after reboot or GRES, the show chassis routing engine command shows incorrect output. The 5 seconds, 1, 5 and 10 minute CPU average utilization is not shown in the output. Its a display issue and there is no functional impact due to this issue. [PR1812514](#)
- Unsupported PEM/PSU is shown as online (green)in the MIST Dashboard and the output of show chassis environment for that PSU shows the status as present/OK. No functional impact. [PR1814463](#)
- In a specific configuration change after NSSU, that is delete and add sequence of link aggregation bundles (LAG) done through load baseline configuration and re-apply original configuration. OSPF session might get stuck in EXSTART state. [PR1817034](#)
- Time Domain Reflectometry (TDR) support for detecting cable breaks and shorts aborts intermittently on some random ports. [PR1820086](#)
- An FXPC core might be generated when an offline and online activity is performed on a 1x100GE Uplink module. The system resumes in normal fashion after the core. [PR1823097](#)
- On the EX4400-48MXP/48XP devices with 1x100G or 4x25G Uplink Module(ULM) in PIC slot 2, when we perform a PIC offline/online operation, we might see messages related to CPU hog by the threads CMQFX or Task ACQUIRE_FP_LOCK. These will be seen only during the operation and does not affect the offline or online operation of the PIC. [PR1823394](#)
- While performing a 4x25g channelization configuration on the 1x100GE PIC, following error logs are printed multiple times momentarily. They are transient log messages. < . . . > **qsfp_tk_cdr_control: qsfp-0/2/0 channelization not yet supported.**< . . . > These message will appear while applying the below configuration #set chassis fpc 0 pic 2 port 0 channel-speed 25g #commit. These are harmless logs and can be ignored. There is no functional impact due to these logs. [PR1823743](#)
- In EX platforms, some EVPN VxLAN T5 routes will not pass traffic after a routing-engine switchover (GRES). [PR1823764](#)
- EX4400 series: When an offline and an online command is issued for a PIC 2 installed with a 1x100GE Uplink module configured for Virtual-chassis operation, the link might not come back to operational state and remains down. [PR1826147](#)
- Virtual chassis auto-conversion is attempted for 4x25G Uplink modules when inserted with 10G-SFP-T transceivers. Auto-conversion fails and links remain down. [PR1826410](#)

- On an EX4400 device with 4x25G Uplink module configured in 1GE or 25G speed, peer side of an interface with 10GBASE-T transceiver may remain up even when the IFL(xe-x/2/y) is not created. For this to happen, a speed mismatched configuration is needed, where a 1G speed or a 25G speed is configured on the PIC 2. [PR1831409](#)

Infrastructure

- When storage on hardware is full, a panic resulting in vmcore and is dumping the core file in /var/ crash leads to memory storage in negative, and potentially results in a truncated vmcore. [PR1796186](#)

Interface and Chassis

- DCD_CONFIG_WRITE_FAILED: IFL me0.0 configuration write failed for an IFL ADD: File exists - On some rare occurrence, the aforementioned message might get filled up in logs after an upgrade to 24.2R1-S1 image. [PR1827981](#)

Platform and Infrastructure

- On EX4300 or EX4300-VC, removal of a Physical Interface Card (PIC), or if the software fails to detect a PIC that is installed, it can cause a crash in the pfex process. This crash can lead to high CPU usage and potentially disrupt network traffic. [PR1779410](#)
- On EX4300 Platforms, Packet Forwarding Engine (PFE) crash will be seen due to an unexpected switchover after committing interface configuration. [PR1785058](#)
- A vmcore and a ksyncd core might be generated during junos image upgrade. [PR1827102](#)
- When a EX4400-48MP/ EX4400-48MXP is rebooted, sometimes the multirate gigabit ethernet (mge-) interface fails to receive traffic after booting up. However, the interface remains operational at the time of issue. The link is up. [PR1827455](#)
- In a EX4400 Virtual Chassis with multiple member stacked together with PIC 2 as VC Port, a reboot or a routing-engine switch-over might result in a VC Port going down specifically operating on a 4x25GE Uplink module. The aforementioned problem might be seen intermittently. [PR1829037](#)

Resolved Issues

IN THIS SECTION

- [General Routing | 41](#)
- [Interfaces and Chassis | 42](#)
- [Layer 2 Ethernet Services | 42](#)

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

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

General Routing

- TSC_DEADLINE disabled due to Errata; please update microcode to version: 0x3a (or later) seen upon upgrade to 21.4 [PR1608045](#)
- New option to allow operator to configure a power source alarm if power is not received on PD ports. [PR1722976](#)
- EX-hardening: EX4400: set chassis config-button no-clear is not working. [PR1758042](#)
- Peer device ports connected to Gigabit Physical ports of EX4100 transition to up state momentarily during reboot of EX4100. [PR1775479](#)
- On MX and EX platform replacing the line-cards may trigger FPC to be offlined due to unreachable destinations. [PR1777534](#)
- Error is shown on system when pvidb variable is accessed. [PR1781317](#)
- MPC line card crashes while ISSU to Junos OS Release 24.2 or later, displays "ISSU PREPARE TIMEOUT" error. [PR1785960](#)
- Interface configured with BPDU-disable goes down during VC mastership switchover. [PR1787892](#)
- Traffic loss after PIC restart if the packet has a VLAN tag of 4095. [PR1788573](#)
- Watchdog SPI transaction is causing the interface flap in the system. [PR1789272](#)

- On EX2300/EX3400 series SFP-SX interface is not come up due to auto-negotiation failure. [PR1789617](#)
- The l2ald process will crash, with rapid configuration changes followed by rpd and l2ald restart process. [PR1790064](#)
- The access port is dropping a VLAN-tagged packet of which the interface is a VLAN-member. [PR1790316](#)
- Warning message 'Too many VLAN-IDs on untagged interface' is seen when more than 1025 vlans on the same LAG interface are configured. [PR1791053](#)
- With any config change or interface up/down with MACsec protocol configured with or without Dot1x, dot1xd process core is observed in the device. [PR1792507](#)
- Dot1x process crash will be seen in the system with "server-timeout" and "server-fail use-cache" configuration. [PR1794778](#)
- On all Junos EX platforms rewrite rules does not work properly when multiple interfaces are configured. [PR1795545](#)
- Cos rewrite rules does not work properly when input/output-vlan-map swap are configured. [PR1795807](#)
- DHCP IP assignment will fail on VoIP phone connected to a VXLAN access port. [PR1797422](#)
- Intermittent alarms related to fan overspeed value can be observed on EX4100 platform. [PR1797727](#)
- [EX2300]"Ethernet Link Down" would not be generated when me0 was down. [PR1799093](#)
- CPU usage gets spiked for eventd due to flooding of pfe_khms_spurious_wakeup log. [PR1801535](#)
- set chassis config-button no-clear support not added on EX4100. [PR1802614](#)

Interfaces and Chassis

- The ifinfo process crash is seen on Junos platforms. [PR1786555](#)

Layer 2 Ethernet Services

- DHCP clients are not receiving IP address. [PR1776451](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 43

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

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

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

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

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



NOTE: The sixty months of support for EEOL releases is introduced in Junos OS 23.2 release and is available for all later releases. For releases prior to 23.2, the support for EEOL releases continues to be thirty six 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 2: 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)	60 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 | 45](#)
- [Known Limitations | 45](#)
- [Open Issues | 45](#)
- [Resolved Issues | 46](#)
- [Migration, Upgrade, and Downgrade Instructions | 46](#)

What's New

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

What's Changed

IN THIS SECTION

- [EVPN | 45](#)

Learn about what changed in this release for JRR Series Route Reflectors.

EVPN

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the `[edit protocols evpn]` hierarchy level. In most use cases, you don't need to change the default limit. If you want to change the default limit, we recommend that you don't set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

[See [mac-ip-limit](#).]

Known Limitations

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

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

Open Issues

There are no known issues in hardware or software in this release for JRR Series Route Reflectors.

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

Resolved Issues

There are no resolved issues in this release for JRR Series Route Reflectors.

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

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 46

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

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

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

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

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



NOTE: The sixty months of support for EEOL releases is introduced in Junos OS 23.2 release and is available for all later releases. For releases prior to 23.2, the support for EEOL releases continues to be thirty six 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)	60 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 MX Series

IN THIS SECTION

- [What's New | 48](#)
- [What's Changed | 95](#)
- [Known Limitations | 101](#)
- [Open Issues | 103](#)
- [Resolved Issues | 111](#)
- [Migration, Upgrade, and Downgrade Instructions | 121](#)

What's New

IN THIS SECTION

- [Hardware | 49](#)
- [Authentication and Access Control | 65](#)
- [Chassis | 65](#)
- [Class of Service | 66](#)
- [Dynamic Host Configuration Protocol | 67](#)
- [EVPN | 67](#)
- [Forwarding Options | 69](#)
- [High Availability | 69](#)
- [Interfaces | 69](#)
- [Juniper Extension Toolkit \(JET\) | 70](#)
- [Junos OS API and Scripting | 70](#)
- [Junos Telemetry Interface | 71](#)
- [MPLS | 75](#)
- [Network Management and Monitoring | 80](#)
- [Precision Time Protocol \(PTP\) | 83](#)
- [Public Key Infrastructure \(PKI\) | 84](#)
- [Routing Policy and Firewall Filters | 84](#)
- [Routing Protocols | 85](#)
- [Securing GTP and SCTP Traffic | 89](#)
- [Serviceability | 89](#)
- [Services Applications | 90](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 91](#)
- [Software Installation and Upgrade | 92](#)
- [Subscriber Management and Services | 94](#)
- [Additional Features | 94](#)

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

To view features supported on the MX Series platforms, view the Feature Explorer using the following links. To see which features are supported in Junos OS Release 24.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- [MX150](#)
- [MX204](#)
- [MX240](#)
- [MX304](#)
- [MX480](#)
- [MX960](#)
- [MX2008](#)
- [MX2010](#)
- [MX2020](#)
- [MX10003](#)
- [MX10004](#)
- [MX10008](#)
- [MX10016](#)
- [vMX](#)

Hardware

- **New HVAC/HVDC Power Supply Unit (PSU) (MX10004, and MX10008)**—Starting in Junos OS Evolved Release 24.2R1, we introduce a new high-capacity HVAC/HVDC PSU (JNP10K-PWR-AC3H). The JNP10K-PWR-AC3H PSU is designed to support high voltage AC or high voltage DC systems in a 15 A or 20 A power mode. The PSU takes four single-phase HVAC (180-305 VAC) or HVDC (190 - 410VDC) inputs at either 20 A or 15 A and provides a DC output of 12.3V. The power supply contains five DIP switches on the faceplate to let you configure the power supply for high-power (20 A) or low-power (15 A) input mode.

[See [MX10004 Power System](#), and [MX10008 Power System](#).]

- **New DC Power Supply Unit (PSU) (MX10004, and MX10008)**—Starting in Junos OS Release 24.2R1, we introduce a new high-capacity DC PSU (JNP10K-PWR-DC3). The JNP10K-PWR-DC3 PSU is designed to support four power supplies in a single housing that accepts either 60 A or 80 A power from four input power feeds. The power supply contains five DIP switches on the faceplate to let you configure the power supply for high-power (80 A) or low-power (60 A) input mode.

[See [MX10004 Power System](#), and [MX10008 Power System](#).]

- **New MIC for MPC2E-3D and MPC3E-3D line cards**—Starting in Junos OS Release 24.2R1, the MPC2E-3D and MPC3E-3D line cards support MIC-3D-10GbE-SFP-E Modular Interface Card (MIC). This MIC offers crucial 1 G/10 GbE port compatibility for MPC2E-NG and MPC3E-NG supported line cards on the MX240, MX480, MX960, MX2010, and MX2020 chassis. In addition to port compatibility, this MIC offers crucial capabilities, including Media Access Control Security (MACsec) and Precision Time Protocol (PTP) (Class B) support, catering to timing applications and meeting the requirements of federal and service provider customers. This MIC ensures seamless integration and compatibility across various deployment scenarios on the MX Series platforms.

Table 4: Features Supported for MIC-3D-10GE-SFP-E

Features	Description
Chassis	<ul style="list-style-type: none"> • The 10x10GbE SFPP and 10x1GbE SFP MIC supports the MPC2E-NG and MPC3E-NG line cards on the MX240, MX480, MX960, MX2010, and MX2020 routers. This MIC has 10 ports that support 1-Gbps small form-factor pluggable (SFP) and 10-Gbps small form-factor pluggable plus (SFP+) transceivers along with PTP and MACsec capabilities. <p>[See pic-mode and number-of-ports.]</p>
Hardware	<ul style="list-style-type: none"> • We support a new MIC, MIC-3D-10GbE-SFP-E, for the MPC2E-3D-NG, MPC2E-3D-NG-Q, MPC3E-3DNG, and MPC3E-3D-NG-Q line cards. The MIC has ten 10GbE ports that support SFP and SFP+ transceivers. The ports provide MACsec support. <p>[See MICs Supported by MX Series Routers.]</p>
High availability and resiliency	<ul style="list-style-type: none"> • Support for MIC (MIC-3D-10GbE-SFP-E) resiliency in MX Series devices.

Table 4: Features Supported for MIC-3D-10GE-SFP-E *(Continued)*

Features	Description
Interfaces	<ul style="list-style-type: none"> • Support for MIC with 1GbE SFP or 10GbE SFP+ ports along with PTP and MACsec capabilities on the line cards. [See Port Speed on MX Routers] • Supported transceivers, optical interfaces, and DAC cables—Select your product in the Hardware Compatibility Tool to view the supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available. • Oversubscription and preclassification support for MIC MIC-3D-10GbE-SFP-E on MX Series devices. [See Oversubscription.]

Table 4: Features Supported for MIC-3D-10GE-SFP-E (*Continued*)

Features	Description
MACsec	<ul style="list-style-type: none"> • Support for Media Access Control Security (MACsec) on physical and virtual interfaces with GCM-AES-128, GCM-AES-256, GCM-AES-XPN-128, and GCM-AES-XPN-256 encryption. Both physical and virtual interfaces support static connectivity association key (CAK) mode. Only physical interfaces support dynamic CAK mode, preshared key (PSK) hitless rollover keychain, and aggregated Ethernet. <p>[See Configuring MACsec.]</p> <ul style="list-style-type: none"> • Precision Time Protocol with Media Access Control Security encryption (MIC-3D-10GbE-SFP-E) enables the simultaneous support of Precision Time Protocol (PTP) and Media Access Control Security (MACsec) encryption on a single port. <p>The following limitations apply:</p> <ul style="list-style-type: none"> • The maximum number of MACsec-enabled logical interfaces (IFL) is 200 per system. • The maximum number of MACsec-enabled ports with physical interfaces (IFDs) and IFLs where MACsec and PTP are enabled together on different ports is 200 per system. • The maximum number of IFLs that can be supported on both 1G and 10G ports is 128. • PTP in clear text mode is not supported.

Table 4: Features Supported for MIC-3D-10GE-SFP-E (*Continued*)

Features	Description
Timing	<ul style="list-style-type: none"> • Synchronous Ethernet with G.8262 standard support on MIC-3D-10GbE-SFP-E. We support Synchronous Ethernet with G.8262 in compliance with the following International Telecommunication Union Telecommunication Standardization (ITU-T) standard to facilitate the transference of clock signals over the Ethernet physical layer. Synchronous Ethernet (G.8262). Timing and synchronization aspects in packet networks. Specifies timing characteristics of synchronous Ethernet equipment clock (EEC). [See Synchronous Ethernet.] • Precision Time Protocol with G.8275.1 standard support on MIC-3D-10GbE-SFP-E. We support Precision Time Protocol with G.8275.1 in compliance with the following International Telecommunication Union Telecommunication Standardization (ITU-T) standards to facilitate distribution of precise time and frequency over packet-switched Ethernet networks. <ul style="list-style-type: none"> • G.8275.1—PTP profile for phase and time (full timing support) • G.8275.1—PTP profile for phase and time over link aggregation group (LAG) [See Precision Time Protocol.]

- **New MX10K-LC4800 line card (JNP10K-LC4800)**—Starting in Junos OS Release 24.2R1, we introduce the MX10K-LC4800 line card that supports three PICs per line card. PIC 0 and PIC 1 have 12 SFP56-DD ports that support 100 Gbps and two QSFP56-DD ports that support 400 Gbps. PIC 2 has 16 SFP56-DD ports that support 100 Gbps. The MX10K-LC4800 line card can deliver up to 4.8 Tbps per-slot bandwidth. The line card interoperates with MX10008 SFB2 and MX10004 SFB2.

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008

Feature	Description
Chassis	<p>Support for the following hardware components, platform features, and fabric functionalities for the new MX10K-LC4800 line card in MX10004 and MX10008 routers:</p> <ul style="list-style-type: none"> Power supply modules (PSMs)— JNP10K-PWR-AC2, JNP10K-PWR-AC3, and JNP10K-PWR-DC2 <p>Use the chassis ambient-temperature (25C 40C 55C) and chassis pfe <pfe Number> power <on off> commands for power management.</p> <p>[See ambient-temperature, and ambient-temperature.]</p> <ul style="list-style-type: none"> Fan trays—JNP10004-FAN2 and JNP10008-FAN2 Interoperability with MX10K-LC480, MX10K-LC2101, and MX10K-LC9600 line cards on MX10004 and MX10008 routers. <p>[See Fabric Plane Management.]</p> <p>If compatible components are not present, SFB2 and the line card continue to remain in offline state and an error message is displayed in the show chassis fpc and show chassis sfb command outputs.</p> <p>[See show chassis fpc and show chassis sfb.]</p> <ul style="list-style-type: none"> Fabric plane management—SFB2 with 6 Packet Forwarding Engines per slot and 12 fabric planes, with fabric fault handling and fabric hardening. <p>[See Fabric Plane Management.]</p> <ul style="list-style-type: none"> Platform resiliency for the following hardware components: <ul style="list-style-type: none"> CPU Field replaceable units (FRUs) Memory

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 (Continued)

Feature	Description
	<ul style="list-style-type: none"> • Management Ethernet ports • Field-programmable gate array (FPGA) board • Optics panel • Power supply module (PSM) • Fan tray <p>If a failure is detected on a hardware component, Junos OS:</p> <ul style="list-style-type: none"> • Logs the message to give clear indication of failure details, including timestamp, module name, and component name. • Raises and clears alarms, if applicable. • Makes the LED glow to indicate FRU fault, if an LED is present. • Performs local action, such as self-healing and taking the component out of service. <p>[See Chassis-Level User Guide.]</p> <ul style="list-style-type: none"> • Junos OS environment monitoring (EM) policy extended to include optics temperature sensors. It includes the following features: <ul style="list-style-type: none"> • The Optics EM policy incorporates periodically polled temperature readings of optical modules in the system to automatically manage the fan speed. • 100GbE and 400GbE optics firmware automatically triggers optics shutdown when the high-temperature threshold is breached. • EM policy is enabled by default on all 100GbE and 400GbE optics interfaces, except for loopback optics and direct attach copper (DAC) cables.

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 (Continued)

Feature	Description
	<p>You can use the <code>set chassis fpc fpc_slot pic pic_slot port port_no no-temperature-monitoring</code> command to explicitly disable the EM policy on specific WAN ports. Use the <code>show chassis environment</code> command to view the optics temperature. [See temperature-sensor.]</p>
CoS	<ul style="list-style-type: none"> • Support for CoS configuration. <p>[See CoS Features and Limitations on MX Series Routers and Hierarchical Class of Service for Subscriber Management Overview.]</p> <ul style="list-style-type: none"> • Support for CoS on 1GbE ports.
DHCP	<ul style="list-style-type: none"> • Support for DHCP functionality (DHCPv4/v6 server, relay, and client feature). [See DHCP Relay Agent.]
Hardware	<ul style="list-style-type: none"> • New MX10K-LC4800 line card (model number: JNP10K-LC4800) on MX10004 and MX10008—We introduce the MX10K-LC4800. MX10K-LC4800 is a fixed-configuration line card that can deliver a bandwidth of up to 4.8 terabit per second (Tbps). The line card has forty 100GbE SFP56-DD ports and four 400GbE QSFP56-DD ports. MX10K-LC4800 interoperates with existing MX Series line cards, such as MX10K-LC9600, MX10K-LC2101, and MX10K-LC480 and interfaces with JNP10004-SF2 (in MX10004) and JNP10008-SF2 (in MX10008) SFBs.
High availability (HA)	<ul style="list-style-type: none"> • Resiliency support for Packet Forwarding Engine and SFB 2. [See show system errors active.]

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 *(Continued)*

Feature	Description
Interfaces	<ul style="list-style-type: none"> • Supports transceivers, optical interfaces, and direct attach copper (DAC) cables on MX10004 and MX 10008. [See HCT and optics-options.] • Interface support on MX10K-LC4800 line card (MX10004 and MX10008)—We introduce the MX10K-LC4800. MX10K-LC4800 is a fixed-configuration line card that can deliver a bandwidth of up to 4.8 Tbps. The line card has forty 100GbE SFP56-DD ports and four 400GbE QSFP56-DD ports. A PIC port in MX10K-LC4800 can support multiple port speeds. • Bit error rate (BER) monitoring and pseudorandom binary sequence (PRBS) support. [See Verifying Link and Transceivers using Pseudo Random Binary Sequence (PRBS) Test and Configuring BERT Testing.]

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 *(Continued)*

Feature	Description
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> • We support Junos telemetry interface (JTI) sensors in the following areas: <ul style="list-style-type: none"> • Chassis management error (cmerror) configuration and counters • Fabric, optical, and Flexible PIC Concentrator (FPC) environment statistics • Platform, interface, and alarm statistics • Transceiver statistics • Segment routing-traffic engineering (SR-TE) for colored telemetry statistics <p>JTI sensor support is unavailable for these actions:</p> <ul style="list-style-type: none"> • Enabling zero suppression in the Packet Forwarding Engine infra for ULC-based line cards • Enabling TARGET-DEFINED support (aft-telemetry library) • Configuring INITIAL-SYNC from the Packet Forwarding Engine Infra for Advanced Forwarding Toolkit (AFT)-based and ULC-based line cards <ul style="list-style-type: none"> • Junos YANG Data Model Explorer--> Telemetry Support. See [https://apps.juniper.net/ydm-explorer/.]

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 *(Continued)*

Feature	Description
L2 features	<ul style="list-style-type: none"> • Layer 2 (L2) learning, trunk port, bridging, integrated routing and bridging (IRB), QinQ, VLAN handling, MAC accounting, LACP, LLDP, xSTP, and Ethernet ring protection switching (ERPS). [See Understanding Q-in-Q Tunneling and VLAN Translation, Understanding Layer 2 Bridge Domains on MX Series, Understanding Layer 2 Learning and Forwarding, and Introduction to OAM Connectivity Fault Management (CFM).] • VRRP [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).] • Support for EVPN-MPLS unicast and multicast forwarding features. [See EVPN User Guide.] • Support for EVPN-VXLAN unicast features. [See Understanding Programmable Flexible VXLAN Tunnels.] • Support for seamless MPLS L2 feature pseudowire headend termination (PWHT). [See Layer 2 VPNs and VPLS Feature Guide for Routing Devices, and Pseudowire Subscriber Logical Interfaces Overview.]

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 *(Continued)*

Feature	Description
L3 features	<ul style="list-style-type: none"> • Support for Layer 3 (L3) features and IGP (OSPF, IS-IS, RIP, and ECMP) for IPv4 and IPv6. [See Understanding OSPF Configurations and BGP Overview.] • Support for forwarding features including IGP, Static Routing, BGP, BGP-PIC, and programmable routing protocol process (rpd). [For OSPF and BGP, see Understanding OSPF Configurations and BGP Overview.] • Support for Bidirectional Forwarding Detection (BFD). [See Understanding BFD for Static Routes for Faster Network Failure Detection.] and Bidirectional Forwarding Detection (BFD).] • SR-TE statistics for uncolored SR-TE policies on JTI. [For SR-TE telemetry, see Understanding OpenConfig and gRPC on Junos Telemetry Interface and Junos YANG Data Model Explorer.] • Support for NGMVPN, GTM, Rosen MVPN, P2MP - RSVP-TE/ mLDP, and MoFRR inband multicast LDP signaling. [See MPLS Overview, Multicast Overview, and Understanding Next-Generation MVPN Control Plane.] • Support for redundant logical tunnel (RLT) interfaces and pseudowire subscriber interfaces using either a logical tunnel or RLT interfaces as an anchor point. [See Redundant Logical Tunnels Overview and Pseudowire Subscriber Logical Interfaces Overview.]

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 (Continued)

Feature	Description
Licensing	<ul style="list-style-type: none"> Support for Licensing: Juniper Agile Licensing provides simplified and centralized license administration and deployment. If you exceed the bandwidth capacity license, Junos OS generates periodic alarms indicating that you need bandwidth capacity license. View the bandwidth information of your device using the <code>show system license bandwidth flex-only</code> command. View the features available on your device using the <code>show system license feature-list</code> command. <p>[See System Log Explorer, Software Licenses for MX Series Routers and MPC Service Cards, and Activate Your Licenses.]</p>
MACsec	<ul style="list-style-type: none"> Support for Media Access Control Security (MACsec) in static connectivity association key (CAK) mode with GCM-AES-128, GCM-AES-256, GCM-AES-XPB-128, and GCM-AES-XPB-256 encryption. Only physical interfaces support this feature. <p>[See Configuring MACsec.]</p>
Network management and monitoring	<ul style="list-style-type: none"> Support for port mirroring. <p>[See Configuring Port Mirroring on MX, ACX, and PTX Series Routers.]</p> <ul style="list-style-type: none"> Support for sFlow. [See sFlow Technology Overview.]
Platform and infrastructure	<ul style="list-style-type: none"> Secure boot and common BIOS implementation based on the UEFI standard. <p>[See Junos OS Overview.]</p>
Routing policy and firewall filters	<ul style="list-style-type: none"> Support for distributed denial of service (DDoS). <p>[See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview .]</p>

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 (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) • MAP-T (Mapping of Address and Port using Translation) <p>[See [Inline NAT, Configuring Mapping of Address and Port with Encapsulation (MAP-E), Configuring Mapping of Address and Port with Translation (MAP-T), and Configuring Inline 6rd.]</p> • RPM, including all probe types and hardware timestamping of RPM probe messages in the Packet Forwarding Engine using the hardware-timestamp and one-way-hardware-timestamp configuration statements. <p>[See Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches, hardware-timestamp , one-way-hardware-timestamp, and rpm (Services).]</p> <ul style="list-style-type: none"> • Two-Way Active Measurement Protocol (TWAMP) <ul style="list-style-type: none"> • Managed control type, for IPv4 traffic • Light control type, for IPv4 and IPv6 traffic • Configuring the TWAMP client instance to use si-x/y/z as the destination interface, which enables inline services, is not supported if the router has a JNP10K-LC4800 line card installed in the chassis. <p>[See Understand Two-Way Active Measurement Protocol, twamp.]</p> • Inline monitoring services support, including support for L2 firewall families. [See Inline Monitoring Services Configuration.] • Juniper Resiliency Interface (JRI): IP Flow Information Export (IPFIX) exception reporting support. [See Juniper Resiliency Interface.]

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 (Continued)

Feature	Description
	<ul style="list-style-type: none">• Inline active flow monitoring support. [See Understand Inline Active Flow Monitoring.]• Routing Engine-based traffic sampling support. [See Configuring Traffic Sampling on MX, M and T Series Routers.]• Inline video monitoring support. [See Inline Video Monitoring.]• FlowTapLite support. [See Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs.]

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 *(Continued)*

Feature	Description
Subscriber management and services	<ul style="list-style-type: none"> • Support for subscriber services functionalities on MX10K-LC4800 line card for MX10004 and MX10008 devices, that includes: <ul style="list-style-type: none"> • PPP subscribers with unauthenticated dynamic VLAN • PPP subscribers with unauthenticated dynamic S-VLAN • PPP subscribers with authenticated dynamic VLAN • PPP subscribers with authenticated dynamic S-VLAN • DHCP subscribers with unauthenticated dynamic VLAN • DHCP subscribers with unauthenticated dynamic S-VLAN • DHCP subscribers with authenticated dynamic VLAN • DHCP subscribers with authenticated dynamic S-VLAN • L2TP subscribers (LNS) • L2TP subscribers (LAC) • Pseudowire • CoS Service • Firewall service • LTS accounting (Goal for FRS) • BFD liveness • CoS support for BNG on pseudowire over active/active RLT interface • L2 dynamic overhead adjust for accounting • Scaling and performance

Table 5: Features Supported on MX10K-LC4800 Line Card for MX10004 and MX10008 (Continued)

Feature	Description
	<ul style="list-style-type: none"> RFC 2544-based benchmarking tests support. [See Understanding RFC2544-Based Benchmarking Tests on MX Series Routers.]

- **New Routing Engine RE-S-X6-128G-LT (MX240, MX480, and MX960)**—In Junos OS Release 24.2R1, we introduce a new Routing Engine, the RE-S-X6-128G-LT. This new Routing Engine is an upgrade to the existing Routing Engine RE-S-X6-64G-LT.



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

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

Authentication and Access Control

- **RADSEC support**—Starting in Junos OS release 24.2R1, RADSEC is supported for ACX, EX, MX, and SRX switches. The RADSEC protocol provides secure transport of RADIUS authentication and accounting data across untrusted networks using Transport Layer Security (TLS) over TCP as the transport protocol.

Chassis

- **10x10GbE SFPP and 10x1GbE SFP MIC support (MX240, MX480, MX960, MX2010, and MX2020 routers)**—Starting in Junos OS Release 24.2R1, the Modular Interface Cards (MICs) 10x10GbE SFPP and 10x1GbE SFP support the MPC2E-NG and MPC3E-NG line cards on the MX240, MX480, MX960, MX2010, and MX2020 routers. This new MIC has 10 ports that support 1-Gbps SFP and 10-Gbps SFP+ speeds along with Precision Time Protocol (PTP) and Media Access Control Security (MACsec) capabilities.
[See [pic-mode \(MX Series\)](#) and [number-of-ports.](#)]
- **Support for Gen3 FT and FTC SKUs (MX-Series)**—Starting in Junos OS Release 24.2R1, support is provided for the new fan trays (FT) (JNP10004 Fan-Tray Gen3 and JNP10008 Fan-Tray Gen3) and fan tray controllers (FTC) (JNP10004 Fan Controller Gen3 and JNP10008 Fan Controller Gen3) SKUs along with resiliency support for MX10004 and MX10008 devices.
- **Optics EM policy support (MX304)**—Starting in Junos OS Release 24.2R1, we have extended the Junos Environment Monitoring (EM) policy to include optics temperature sensors for MX304 routers. It includes the following features:

- The Optics EM policy incorporates periodically polled temperature readings of optical modules in the system to automatically manage the fan speed.
- 100GbE and 400GbE optics firmware automatically triggers optics shutdown when the high-temperature threshold is breached.
- EM policy is enabled by default on all 100GbE and 400GbE optics interfaces, except for loopback optics and direct attach copper (DAC) cables.

You can use the `set chassis fpc fpc_slot pic pic_slot port port_no no-temperature-monitoring` command to explicitly disable the EM policy on specific WAN ports. Use the `show chassis environment` command to view the optics temperature.

[See [temperature-sensor](#).]

- **Source Redundancy and Feed Redundancy support (MX10004 and MX10008)**—N+1 power redundancy is supported on MX10004 and MX10008 routers with the JNP10K-PWR-DC3 power supply modules (PSMs). You can enable either source redundancy or feed redundancy for the PSM.

[See [Power Redundancy for Third-Generation Power Supply Modules](#).]

- **Resiliency support (MX10004 and MX10008)**—We support resiliency for JNP10K-PWR-DC3 power supply modules (PSMs) on MX10004 and MX10008 devices. Resiliency enables the system to monitor component health, alert you of errors, and take appropriate action to restore normal operation based on error severity.

[See [Resiliency](#), [thermal-health-check](#), and [watchdog \(PSM\)](#).]

Class of Service

- **Policy map support for inet6-precedence (MX480, MX960, MX10003, MX10004, MX10008, MX10016, and MX2020)**—Policy map is a packet marking scheme that enables you to define rewrite rules on a per-customer basis (that is, for each customer). The policy map makes it possible to use any packet field to identify a given flow and specify a rewrite value for that flow. Starting in Junos OS Release 24.2R1, MX480, MX960, MX10003, MX10004, MX10008, MX10016, and MX2020 routers support these inet6-precedence code point options for policy maps:

- `inet6-precedence proto-ip`
- `inet6-precedence proto-mpls`

You define the `inet6-precedence` code point options for a policy map at the `[edit class-of-service policy-map policy-map-name]` hierarchy level.

[See [policy-map](#).]

- **Detail and summary slice statistics support on physical interfaces (MX240, MX304, MX480, MX960, MX10003, MX10004, and MX10016)**—Starting in Junos OS Release 24.2R1, the `show interfaces queue`

interface-name slice command includes the detail and summary options, as well as filtering sub-options to display detail and summary statistics for all slices on an interface.

[See [show interfaces queue](#).]

Dynamic Host Configuration Protocol

- **Support for DHCPv6 server address pool obtained for router advertisement (NFX Series)—**

When an NFX router functions as both the DHCPv6 client and server, the prefix information received from the WAN side through DHCP-prefix delegation (PD) needs to be divided. This prefix information is used in multiple LAN side interfaces for router advertisement or for DHCPv6 server address pools. Starting in Junos OS Release 24.2R1, NFX Series routers support consistent division of the received prefix into sub-prefixes by providing new optional configuration for router advertisement or for the DHCP local server.

See [dhcpv6-client](#).

- **Support for deferred negative acknowledgement (deferred-NAK)** Starting in Junos OS 24.2R1 release, the existing JUNOS DHCPV4/V6 Server 'FORCE-RENEW/RECONFIGURE' support now also supports a "deferred-NAK" option, whereby if the DHCP client does not immediately respond to the FORCE-RENEW/RECONFIGURE request (or any of its subsequent limited retries), the subscriber session is left in place with full connectivity and the session gets flagged for "deferred-NAK". This session state is maintained persistently across daemon restarts and GRES/ISSU events.

EVPN

- **Support for single-link targeting on redundant logical tunnel (MX240, MX480, MX960, MX10003, MX1008, and MX10016)—**Starting in Junos OS Release 24.2R1, we support single-link targeting. When you configure single-link targeting, all subscribers using the RLT are terminated when the targeted logical tunnel link goes down.

[See [Logical Tunnel Interfaces](#) and [Redundant Logical Tunnels](#).]

- **Support for minimum active links on RLT (MX250, MX480, MX960, MX10003, MX10008, and MX10016)—**Starting in Junos OS Release 24.2R1, you can configure a minimum number of active links on a redundant logical tunnel (RLT) interface. When the number of up logical tunnel links in the RLT drops below the configured minimum, the RLT goes down. Pseudowire interfaces stacked on the same RLT also go down in this scenario.

[See [Logical Tunnel Interfaces](#), [Redundant Logical Tunnels](#), [Pseudowire Configuration](#), and [Pseudowire Headend Termination \(PWHT\)](#).]

- **RFC7432 compliance for VLAN-based EVPN with an IRB interface (MX960)—**Starting in Junos OS Release 24.2R1, we've added a new CLI statement `advertise-zero-ethernet-tag` for VLAN-based EVPN instances using an IRB interface to provide Layer 3 gateway (L3GW) functionality. You use this configuration to advertise EVPN Type 2 (MAC/IP Advertisement) and EVPN Type 3 (Inclusive

Multicast Ethernet Tag) routes with an Ethernet tag value of 0 for RFC7432 compliance when the instance has a valid `vlan-id`. This statement enables a VLAN-based service to provide L3GW functionality with a `vlan-id` configured and still be RFC7432 compliant for Layer 2 gateway (L2GW) functionality by advertising EVPN routes with an Ethernet Tag ID set to 0 instead of using the `vlan-id`.

[See [advertise-zero-ethernet-tag](#).]

- **EVPN ESI per EVI Instance Support (MX960)**—Starting in Junos OS Release 24.2R1, you can provide Ethernet segment ID (ESI) support at the EVPN instance (EVI) level for EVPN VLAN-based service in active-standby redundancy mode. We've introduced a new statement `per-evi` to support this feature. To enable this feature, configure a unique ESI at the EVI level and configure all access logical interfaces (IFLs) under the EVI with the same ESI value and the `per-evi` statement. When you enable this feature, the designated forwarder (DF) switches to the non-designated forwarder (NDF) role only when all the access IFLs under the EVI go down. You can avoid traffic interruption by preventing a role switch when a single access IFL fails on the EVI.

[See [per-evi](#).]

- **Suppress EVPN Type 5 host routes from DCI to DC (EX4400-24MP, EX4400-48F, MX304, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can suppress EVPN Type 5 host route advertisements that re-originate from the data center interconnect (DCI) to the local DC. You can achieve better scaling and performance on leaf devices with this feature.

[See [suppress-host-routes-from-dci-to-dc](#).]

- **Non-revertive preference-based DF election in EVPN-MPLS networks (EX4400-24P, MX960, and vMX)**—Starting in Junos OS Release 24.2R1, you can configure non-revertive preference-based designated forwarder (DF) election for an Ethernet segment identifier (ESI) in an Ethernet VPN—MPLS (EVPN-MPLS) network. By default, preference-based DF election for an Ethernet segment identifier (ESI) is revertive, which means:
 - If the EVPN provider edge (PE) device currently in the DF role goes down, the next preferred PE device becomes the new DF.
 - When the old DF comes back up, the DF role reverts to the old DF.

Changing the current DF role for an ESI frequently can impact traffic flow. To avoid revertive DF role changes, you can now set the `non-revertive` option at the `[edit interfaces name esi df-election-type preference]` hierarchy level.

We also provide new options you can configure to load-balance DF election per EVPN instance (EVI) or per ESI based on the lowest configured preference value or the highest configured preference value, as follows:

- At the EVI level—Use the designated-forwarder-preference-least option or the designated-forwarder-preference-highest option at the [edit routing-instances *evpn-instance-name* protocols evpn] hierarchy level.
- At the ESI level—Use the least option at the [edit interfaces *interface-name* esi df-election-type preference] or [edit protocols evpn interconnect esi df-election-type preference] hierarchy level.

[See [df-election-type](#) and [evpn](#)].

- [See [dot1x](#).]
- **Support for EVPN-VPWS over SRv6 with micro-SID** (MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, and MX10008)—Starting in Junos OS Release 24.2R1, we support Segment Routing for IPv6 (SRv6) underlay feature with micro segment identifier (micro-SIDs) over Ethernet VPN–virtual private wireless service (EVPN-VPWS)..



NOTE: The MX series of routers support a single segment identifier with this feature.

[See [Configuring Micro-SIDs in EVPN-VPWS](#).]

Forwarding Options

- **Support for consistent-hash over static routes** (MX240, MX304, MX480, MX960, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)—Starting in Junos Release OS 24.2R1, you can enable symmetry with consistent-hash using the `symmetric-consistent-hash` configuration statement configured in conjunction with `source-ip` only load-balance in one direction and `destination-ip` only in reverse direction.
- [See [Actions in Routing Policy Terms](#).]

High Availability

- **Multihop BFD echo-lite sessions for distributed and inline modes (MX960 and MX2010)**—Starting in Junos OS Release 24.2R1, programmable RPD (PRPD) BFD APIs support multihop BFD echo-lite sessions in distributed and inline modes. Multihop BFD echo-lite sessions are used to probe the liveness of peer devices that are multiple hops away, even if the peer device does not support BFD.

[See [show bfd session extensive](#).]

Interfaces

- **TLB: Introducing OR function between TLS and SSL probes (MX240, MX480, and MX960)**—Starting in Junos OS Release 24.2R1, the system uses the *OR* mechanism instead of *AND* to determine the status of the real server when TLS and SSL are configured in the same group. That is, the real server is marked as up if any one of the probes is successful.

When the SSL probing version is provided, the system uses that version to probe. When the SSL version is not specified, the behavior changes to Fallback from version SSLv3 to SSLv2. The probe starts with SSLv3. If the SSLv3 probe fails, the system probes for SSLv2.

[See [Traffic Load Balancer](#).]

- **Support for logical interfaces under a demux interface (MX240, MX480, MX960, MX10004, and MX10008)**—Starting in Junos OS Release 24.2R1, you can configure up to 32,000 logical interfaces (IFLs) under one aggregated Ethernet (ae) interface device. The ae interface must be configured under the demux0 interface to support 32,000 IFL's. If the ae interface is configured under any other interface or device, a 16,000 IFL limit applies. This allows you to configure your device to support twice as many subscriber interfaces over the demux0 interface.

[See [Subscriber Interfaces and Demultiplexing Overview](#), [Dynamic Demultiplexing Interfaces](#), and [Demultiplexing Interface Overview](#).]

Juniper Extension Toolkit (JET)

- **JET Firewall Filter API to configure purge timer and support filter match condition (MX204, MX240, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—Starting in Junos OS Release 24.2R1, use the JET Firewall Service API to configure the purge timer for JET clients. With the default configuration, a JET client purges its configuration if it fails to connect within 30 seconds. When the connection is successful, the client must reconfigure. To avoid delays due to reconfiguration, you can extend the purge timer beyond 30 seconds.

The JET Firewall Filter API also supports firewall filter match with GPRS tunneling protocol (GTP). You can configure this match on the interface to better filter GTP traffic. Note that a GTP filter match works only if you configure the appropriate port match condition for UDP traffic.

[See [Overview of JET APIs](#).]

- **Persistence for PRPD routes (MX10003)**—Starting in Junos OS Release 24.2R1, you can enable persistence for up to 500,000 programmable RPD (PRPD) routes using JET APIs. Persistent routes are maintained across restarts and switchovers, providing reliability for important routes. To enable persistence, set the persist variable to 1 in the RouteAdd, RouteModify, or RouteUpdate remote procedure call (RPC) under the RIB Service API. This feature is only supported for flexible VXLAN routes.

If a persistent PRPD route is configured with a backup route, use the delay-route-advertisements minimum-delay routing-uptime configuration statement to delay BGP route advertisements so the PRPD route can be reprogrammed after a restart.

[See [Overview of JET APIs](#).]

Junos OS API and Scripting

- **Scapy Python library support (MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, and MX10008)**—Starting in Junos OS Release 24.2R1, Junos OS

supports a limited version of the Scapy Python library. Scapy is an interactive packet manipulation library that enables you to create, send, receive, and dissect network packets. You can use Scapy to gather information or troubleshoot issues in your network.

[See [Overview of Python Modules on Devices Running Junos OS](#).]

Junos Telemetry Interface

- **OpenConfig configuration and sensor support for ZR and ZR+ optical transceivers (MX2020)**—Junos OS Release 24.2R1 supports the OpenConfig configuration and data streaming of ZR and ZR+ optics. You can create a subscription in INITIAL_SYNC or TARGET_DEFINED mode using Juniper proprietary Remote Procedure Call (gRPC) service or gRPC Network Management Interface (gNMI). Use the OpenConfig command `/components/component/transceiver/config/fec-mode` for configuration. Use these resource paths in a subscription to stream data:

- `/components/component/transceiver/state/`
- `/components/component/transceiver/physical-channels/`
- `/components/component/optical-channel/state/`

This feature is based on data models `openconfig-terminal-device.yang` (version 1.8.0), `openconfig-platform-transceiver.yang` (version 0.8.0), and `openconfig-transport-types.yang` (version 0.14.0).

[For the optics configuration, see [Mapping OpenConfig Interface Commands to Junos Configuration](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **OpenConfig MAC address and MAC address and IP path sensor support (ACX710, ACX5448, ACX5448-M, ACX5448-D, EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-48MP, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-24T, EX4100-F-12P, EX4100-F-48T, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vMX, QFX10002-60C, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, QFX10002, QFX10008, and QFX10016)**—Junos OS Release 24.2R1 supports the streaming of telemetry MAC address and MAC address and IP path data from the forwarding database to a collector using the OpenConfig resource path `/network-instances/network-instance/fdb/`. This feature is based on data models `openconfig-network-instance.yang` (version 1.2.0) and `openconfig-network-instance-l2.yang` (version 1.2.0).

[See [Junos YANG Data Model Explorer](#).]

- **Support for QoS classifier, rewrite, and scheduler OpenConfig configurations and state sensor support (MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vMX)**—Junos OS Release 24.2R1 supports the QoS classifier, rewrite and scheduler OpenConfig configurations and state sensors using Junos telemetry interface

(JTI). This feature includes support for term ID, behavior aggregate (BA) classifiers for IPv4 and IPv6, and Differentiated Services Code Point (DSCP). The feature also supports term ID for rewrites and the sequence ID and burst size for scheduler maps. We support these OpenConfig configurations:

- `/qos/classifiers/classifier/terms/term/config`
- `/qos/classifiers/classifier/terms/term/config/id`
- `/qos/classifiers/classifier/terms/term/conditions/ipv4/config/dscp-set`
- `/qos/classifiers/classifier/terms/term/conditions/ipv6/config/dscp-set`
- `/qos/scheduler-policies/scheduler-policy/schedulers/scheduler/config/sequence`
- `/qos/scheduler-policies/scheduler-policy/schedulers/scheduler/two-rate-three-color/config/be`

We support these state sensors:

- `/qos/classifiers/classifier/terms/term/state`
- `/qos/classifiers/classifier/terms/term/state/id`
- `/qos/classifiers/classifier/terms/term/conditions/ipv4/state/dscp-set`
- `/qos/classifiers/classifier/terms/term/conditions/ipv4state/dscp-set`
- `/qos/scheduler-policies/scheduler-policy/schedulers/scheduler/state/sequence`
- `/qos/scheduler-policies/scheduler-policy/schedulers/scheduler/two-rate-three-color/state/be`

[For QoS OpenConfig configurations, see [Mapping OpenConfig QoS Commands to Junos Configuration](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **IP oc-aft supports periodic streaming of selected prefixes (MX2020)**—Junos OS Release 24.2R1 supports periodic streaming of specific prefixes under the IP OpenConfig Abstract Forwarding Table (oc-aft) sensor child path `/network-instances/network-instance/afts/`. To enable prefix filtering on the target (source) device, include the prefix statement at the `[edit fib-streaming prefix-list table table-name family family-name]` hierarchy level. When you enable this feature, only interfaces with the required prefixes and their corresponding next-hop and next-hop group containers are exported to the oc-aft collector. Reducing the set of interfaces to only the ones of interest to the collector decreases the overall CPU and resource usage on Routing Engines, Flexible PIC Concentrators (FPCs), and Modular Port Concentrators (MPCs). The recommended periodic interval for streaming resource paths under `/network-instances/network-instance/afts/` is 5 minutes.
[See [Configuring Prefix Filtering](#), `prefix-list`, `show fib-streaming state`, and [Junos YANG Data Model Explorer](#).]
- **OpenConfig compliance and support for SR-TE sensors (MX240, MX304, MX480, MX960, MX2010, MX2020 MX10003, MX10004, MX10008, and MX10016)**—Junos OS Release 24.2R1 supports OpenConfig-compliant Segment Routing–Traffic Engineering (SR-TE) sensors. This feature includes

per segment list data for colored tunnels and new resource paths for existing colored sensors. You can stream data from a device to a collector using native (UDP) sensors, Juniper proprietary Remote Procedure Call (gRPC) service, or gRPC Network Management Interface (gNMI).

[For sensors, see [Junos YANG Data Model Explorer](#).]

- **Platform telemetry sensor support (MX10004 and MX10008)**—Junos OS Release 24.2R1 supports platform sensors that stream telemetry using Juniper proprietary Remote Procedure Call (gRPC) service or gRPC Network Management Interface (gNMI). ON_CHANGE, INITIAL_SYNC, and TARGET_DEFINED subscription modes as well as zero-suppression are supported. The following sensors are supported for all power entry modules (PEMs) and fans:

- `/components/component/properties/property`
- `/components/component/state/temperature`
- `/components/component/state`
- `/components/component/fan/state`
- `/components/component/power-supply/state`

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **Queueing block pipeline sensor support (MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—Junos OS Release 24.2R1 supports queueing block pipeline sensors. Stream data about Packet Forwarding Engine counters at a granular level using Junos telemetry interface (JTI) with native (UDP), Juniper proprietary Remote Procedure Call (gRPC) service, or gRPC Network Management Interface (gNMI). Use these sensors in a subscription to export data:

- `/components/component/integrated-circuit/pipeline-counters/packet/queueing-block/`
- `/components/component/integrated-circuit/pipeline-counters/drop/queueing-block/`

[For sensors, see [Junos YANG Data Model Explorer](#).]

- **QoS OpenConfig and operational state sensor support (MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vMX)**—Junos OS Release 24.2R1 supports QoS configurations and sensors you can use to manage the Junos CoS features drop profile and traffic control profile, as well as stream state data for these features. The OpenConfig data model `openconfig-qos` (version 0.9.1) supports these enhancements.

[For OpenConfig configuration, see [Mapping OpenConfig QoS Commands to Junos Configuration](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **Native inline flow monitoring for IPFIX sensor support (MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vMX)**—Junos

OS Release 24.2R1 supports native streaming of inline flow monitoring for IP Flow Information Export (IPFIX) operational states. To stream data, use the sensor `/state/sampling/flow-monitoring`. [For sensors, see [Junos YANG Data Model Explorer](#).]

- **Additional sensor support for oc-aft streaming of FIB telemetry (MX240, MX960, and MX2020)**—Junos OS Release 24.2R1 extends OpenConfig Abstract Forwarding Table (oc-aft)-based streaming of forwarding information base (FIB) telemetry to support specific leaves under IPv4-unicast, IPv6-unicast, next-hop-groups, and next-hops containers. The support is based on the OpenConfig data model `openconfig-aft`.

You must meet the existing requirements to stream telemetry data and also include the `oc-tlv` statement at the `[edit routing-options forwarding-table]` hierarchy level to enable statistics collection.

[See [forwarding-table](#) and [show fib-streaming](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **OpenConfig configuration and state sensor support for AFI-SAFI policies (MX960, MX10003, MX10004, MX10008, and MX10016)**—Junos OS Release 24.2R1 supports export and import policies under address family indicator (AFI) and subsequent address family identifier (SAFI). OpenConfig support is for IPv4 and IPv6 unicast address families. Junos CLI-configured policies have priority over policies configured with OpenConfig. For example, if you configure a policy using OpenConfig and another policy using Junos OS CLI under the same neighbor, the latter policy takes effect. We support these OpenConfig configurations:

- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi<IPV4/V6_UNICAST>/apply-policy/config/import-policy`
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi<IPV4/V6_UNICAST>/apply-policy/config/export-policy`

We support these state sensors:

- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi<IPV4/V6_UNICAST>/apply-policy/state/import-policy`
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi<IPV4/V6_UNICAST>/apply-policy/state/export-policy`

[For configurations, see [Mapping OpenConfig BGP Commands to Junos Configuration](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **View number of programmable routes for VRF instances and VNIs (MX304 and MX10003)**—Starting in Junos OS Release 24.2R1, you can view the number of programmable routes per virtual routing and forwarding (VRF) instance and virtual network identifier (VNI) using either the Junos OS CLI or using a sensor, stream the data to a collector. To view the number of routes with the CLI, use the `show programmable-rpd clients route-summary display-vni-data` command. To stream the VNI count to a collector using telemetry, include the sensor `state/routing-instances/routing-instance/routing-tables/`

`routing-table/summaries/programmed/clients/client/protocols/protocol/vnis/vni[id='id']` in a subscription.

Use this feature to improve the performance and efficiency of your device. It is more efficient to check the number of routes with telemetry and CLI than to use gRPC API calls.

[For the CLI command, see [show programmable-rpd clients route-summary](#). For sensors, see [Junos YANG Data Model Explorer](#).]

- **OpenConfig configuration and state sensor support for default leaves for IS-IS (MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—Junos OS Release 24.2R1 supports OpenConfig configuration and state sensor support for default leaf values, closing endpoint gaps for the IS-IS protocol. This feature is based on the OpenConfig data model `openconfig-isis.yang` version (version 1.0.0).
[See [Mapping OpenConfig IS-IS Commands to Junos Configuration](#) and [Junos YANG Data Model Explorer](#).]

MPLS

- **Support for GRE tunnels over PWHT interfaces (MX240, MX304, MX480, MX960, MX10003, and MX10004)**—Starting in Junos OS Release 24.2R1, you can set up a GRE tunnel over your pseudowire headend termination (PWHT) interface with existing configuration commands.
[See [Pseudowire Headend Termination \(PWHT\)](#) and [Configuring GRE Tunnel Interfaces](#).]
- **Support for constraint-aware RSVP bypass LSPs (MX204, MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, you can configure RSVP bypass label-switched paths (LSPs) to be aware of and inherit all the path constraints from the primary LSPs. You can also explicitly configure bypass constraints for individual LSPs. With this feature, you can control the MPLS path and prevent bypass LSPs from traversing through a specific geographical area in a global MPLS RSVP network.
[See [Configuring Constraint Aware Bypass LSPs](#).]
- **Support for LDP dual-transport over IPv4 and IPv6 sessions with NSR configuration (MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—Starting in Junos OS Release 24.2R1, you can configure the LDP dual-transport mechanism with NSR support to set up IPv4 and IPv6 sessions. This feature helps to forward IPv4 and IPv6 traffic and to support LDP IPv6 sessions in a routing instance.
[See [Carrier-of-Carrier VPNs](#), [LDP Overview](#), and [LDP Configuration](#).]
- **Support to increase the retry-timer value (MX10004, MX10008, MX10016, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, you can increase the amount of time the ingress router waits between to receive a response. Use the `adaptive-wait-timer` statement at the `[edit protocols mpls]` hierarchy level to configure the minimum period the ingress router must to wait for receiving a path

response message. If the router does not receive any response within the specified time, the wait timer expires and the router terminates the path message and resends the message in the next path maintenance. The default value of the `adaptive-wait-timer` statement is 180 seconds.

To apply the default time to all LSPs, configure the `initial-time` statement at the `[edit protocols mpls adaptive-wait-timer]` hierarchy level.

You can optionally configure the `max-time` statement at the `[edit protocols mpls adaptive-wait-timer]` hierarchy level to set a maximum exponential backoff timer value for the `adaptive-wait-timer` statement. When the first `adaptive-wait-timer` expires, the router continues to retry the path message. In each attempt, the router exponentially backs off the time specified in `adaptive-wait-timer` to get more time to receive the response.

When the exponential backoff time reaches the `max-time` value, the router can no longer back off the waiting time and waits only for the `max-time` period in further attempts. If you configure `initial-time` and `max-time` with the same value, the router waits for the same period during further attempts without any exponential backoff. The default value of the `max-time` statement is 1800 seconds.



NOTE: When the `adaptive-wait-timer` statement is not configured, the router follows the default behavior of waiting for five times the `retry-timer` value that you configure.

You can increase the exponent-base value by configuring the `backoff-multiplier` statement at the `[edit protocols mpls adaptive-wait-timer]` hierarchy level. The default value of the `backoff-multiplier` statement is 2. For example, if the `initial-time` for the `adaptive-wait-timer` is 180 seconds, then with the default `backoff-multiplier` value of 2, the exponentially backed-off values of `adaptive-wait-timer` will be 180 seconds, 360 seconds, 720 seconds, and so on. If you configure `backoff-multiplier` as 3, then the exponentially backed-off values of `adaptive-wait-timer` will be 180 seconds, 540 seconds, 1620 seconds, and so on.

We've enhanced the `show mpls tunnel-manager-statistics` command to additionally display the number of path messages a router sent and the minimum, maximum, and average time a router takes to receive a response. You can see these statistics only at the global level.

- **Support to exclude a list of hops in the RSVP LSP path (MX480, MX960, MX10004, MX10008, MX10016, and vMX)**—Starting in Junos OS Release 24.2R1, you can configure a list of hops to be excluded in the label-switched path (LSP) so that RSVP LSPs avoid those hops and links in the traffic engineering (TE) domain. When an RSVP LSP is signaled in the network, the path message carries the excluded list of hops. When the downstream routers perform loose hop expansion, such as inter-domain LSP or abstract node expansion, the transit routers use the same excluded list of hops that the ingress router uses for path computation. This mechanism enables intermediate routers to avoid the routers included in the excluded hop list. The routers try alternative paths to help with the convergence of LSPs when a complete end-to-end path computation is not possible.

Additionally, ingress routers receive PathErr messages and when computing another path, the routers use a PathErr message sender's address to avoid the link or node that generates an error. Transit routers also need this error avoidance information during retry attempts. RFC4874 defines the exclude hop information and is accepted in RSVP signaling.

To configure LSPs to exclude a list of hops, include the `exclude` statement at the `[edit protocols mpls path path-name next-hop]` hierarchy level. The ingress routers exclude the hops in CSPF computation and are also included in RSVP LSP signaling.

- **Enhancements to RSVP debug and service commands (MX204, MX480, MX960, MX2020, and vMX)**
—Starting in Junos OS Release 24.2R1, we have enhanced the following show commands to help you analyze and debug the following information:
 - History of major events on the label-switched paths (LSPs) and RSVP neighbors that the label-switching router (LSR) maintains
 - Actual time taken for certain events. With this information, you can understand whether certain timer values configured are appropriate or not.

We have enhanced the `show rsvp session extensive` command to display the timeline of the major events that occur for a session and are maintained in a path state block (PSB). The command output also displays message statistics such as Path, Resv, Err sent, and received for a session.

We have enhanced the `show rsvp neighbor` command with the `display level extensive` to display the timeline of major events that take place for an RSVP neighbor.

The `show mpls tunnel-manager-statistics` command has been enhanced to display the minimum, maximum, and average time that clients of an LSP take to relinquish references to an old LSP instance after a make-before-break switchover. These metrics are computed even if the `optimize-adaptive-teardown` statement is not enabled for LSPs.

We have enhanced the `show rsvp statistics` command to display the minimum, maximum, and average time taken for LSPs to be cleaned up by RSVP after the triggering of soft preemption.

We have increased the maximum configurable value of `teardown-timeout` at the `[edit protocols mpls oam bfd-liveness-detection failure-action make-before-break]` hierarchy level from 30 seconds to 65535 seconds.

We have increased the maximum configurable value of `lsp-ping-multiplier` at the `[edit protocols mpls oam lsp-ping-multiplier]` hierarchy level from 1 through 255 (previously 1 through 5).

- **Support for IPv4 static route over IPv6 next-hop (MX204, MX240, MX304, MX480, MX960, MX10003, MX10016, MX2020, QFX5110, and QFX5200)**—Starting in Junos OS Release 24.2R1, you can configure an IPv4 static route over an IPv6 next hop to enable routing of IPv4 packets through the IPv6 next hop. The following IPv4 static route over IPv6 next-hop are supported:
 - IPv4 static route over IPv6 direct next-hop

- IPv4 static route over IPv6 indirect-next-hop
- IPv4 static route over IPv6 next-hop with preference

Use the following configuration to support IPv4 static route over IPv6 next-hop:

```
user@host# set routing-option static route ipv4-address next-hop ipv6-address
```

- **Support for dynamic tunnel for best effort SRv6 tunnels (MX204, MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, you can configure Segment Routing for IPv6 (SRv6) Layer 3 VPN dynamic tunnel over a traditional Layer 3 VPN network.

The following functionalities are supported:

- DT4, DT6, DT46, uDT4, uDT6, uDT46 SIDs.
- Signal SRv6 locator based dynamic tunnel from BGP.
- Resolve BGP route over dynamic tunnel route.
- Resolve BGP route over dynamic tunnel and create transport tunnel composite next hop (TCNH) with BGP, IGP, and static as the underlay. Have single or multiple router next hops.
- Forward policies under dynamic tunnels.
- Propagate DSCP for dynamic tunnel at ingress.
- Display dynamic tunnel (dyn-tunnel) flag information for SRv6 tunnel as part of show route extensive command.
- **Support to distribute the Entropy Label Capability (ELC) in an ISIS network (MX10003, MX10004, MX10008, and MX10016)**—Starting in Junos OS Release 24.2R1, you can distribute ELCs across all the routers in an ISIS network. ELC indicates the capability of a router to interpret Entropy Label Indicator (ELI), remove ELI/EL, and inspect next label. Entropy Readable Label Depth (ERLD) is the number of labels the router is able to read in a label stack and use it for its load balancing function. This can be used in cases of stacked labels (SR-MPLS) to insert ELs at ingress routers based on the different ELC and ERLD of the routers along its path.

You can configure the `entropy-label` statement at the `[edit protocols isis source-packet-routing]` and at the `[edit protocols source-packet-routing source-routing-path <*>]` hierarchy levels to enable this feature. When the `entropy-label` statement is configured, the L-ISIS routes and SRTE for the prefixes are installed with a Entropy Label Indicator (ELI) if the endpoint is entropy label capable. Entropy labels are inserted only at the bottom of the label stack regardless of the ERLD of the routers along the path of the tunnel.

The prefixes with `entropy-label-capability-flag` statement under the `prefix-attribute-flags` in the policy statement is advertised in the router to support entropy label based load balancing.

ELC in ISIS network supports the following functionalities:

- Store ELC in ISIS database.
- Distribute ELC across all the routers participating in the ISIS network.
- Propagate ELC information from ISIS database to TED.
- Reflect ELC capability from TED in the `lsdist` table as a part of the Prefix Attribute flag.
- Reflect ELC capability in the `lsdist` table and TED on the export side.
- Reflect the Prefix Attribute flag in ISIS, TED, and BGP LS on import and export side if `no-load-balance-label-capability` or `load-balance-label-capability` statement is configured or removed.
- Distribute ELC flag across ISIS, TED, and BGP LS if the `entropy-label-capability-flag` statement is added or removed from the policy-statement for the affected prefixes.
- Update L-ISIS routes based on the activation or deactivation of `entropy-label` statement under the `[edit protocols ISIS source-packet-routing]` hierarchy level.
- Update SR-TE routes if prefix of the tunnel endpoint is capable of doing load balancing and `entropy-label` statement is configured or removed.
- Entropy Label Capability flag is preserved when the router propagates the prefix across the ISIS levels.
- Internet, Layer 3 VPN, Layer 3 VPN, and EVPN-based services over SR and SR-TE routes using `entropy-label`.
- Entropy label for both IPv4 and IPv6 prefixes.
- Entropy label for SR-MPLS tunnels with IPv6 endpoint.
- Entropy label for 6PE SRTE tunnels.
- Entropy label capability advertisement for prefixes in different ISIS instance and in multi-topology.
- Entropy label for flex algorithm prefixes.
- Entropy label for source-routing-path-template.
- Entropy label for ping and traceroute to SR-TE tunnel.
- Entropy label for SBFD.

Use the `show isis database`, `show ted database`, and `show route table lsdist.0` commands to view the ELC flag in the Prefix Attribute flags. The `show route` command shows the load balancing capabilities for the L-ISIS and SPRING-TE routes with the entropy label.

The `show spring-traffic-engineering lsp detail` command displays the entropy-label capability of the tunnel only when the `entropy-label` statement is configured for the SR-MPLS in the tunnel or at the instance level.

- **Provision binding SIDs for uncolored SR-TE (SR-MPLS) LSP (MX480 and QFX5200)**—Starting in Junos OS Release 24.2R1, we support provisioning of binding SID for uncolored SR-TE LSP where PCE requests PCC to allocate a binding SID from PCC's label space as follows:
 - PCE requests PCC to allocate a specific binding SID
 - PCE requests PCC to allocate binding SID of PCCs choice

We support the following PCE functionalities:

- PCE requests PCC to allocate binding SID of PCCs choice for delegated LSP.
- PCE requests PCC to allocate binding SID of PCCs choice for PCE-initiated LSP.
- PCE requests PCC to allocate a specific binding SID for delegated LSP.
- PCE requests PCC to allocate a specific binding SID for PCE-initiated LSP.
- Multiple candidate paths with binding SID in a policy.

We now support both 20-bit and 32-bit binding SID provisioned or requested from a PCE controller.

[See [PCEP Configuration](#).]

Network Management and Monitoring

- **Clear LLDP neighbors from an interface with the gRPC Network Operations Interface (gNOI) Layer2 service** (ACX710, ACX5448, ACX5448-M, ACX5448-D, EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200 and QFX5210)—Starting in Junos OS Release 24.2R1, you can execute supported Layer2 service remote procedure calls (RPCs) to perform the equivalent of the `clear lldp neighbors interface interface-name` command.

[See [gNOI Layer 2 Service](#).]

- **CFM CCM support on pseudowire and PWHT service interfaces (MX204, MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, we support connectivity fault management (CFM) check messages (CCM) on PS service interface for active/active mode (with or without targeting), which is part of EVPN.

Pseudowire headend termination (PWHT) allows you to connect an L2 circuit from an access node directly to an L3 service at the service node.

You can configure CCM for down maintenance association end points (MEPs). The CCM MEPs on the PS service interface monitor the Ethernet networks for connectivity faults with continuity check interval of 1 second (1s), 10 seconds (10s), and 100 milliseconds (100ms).

CFM sessions distributed or anchored to the active Packet Forwarding Engine.

The feature support includes:

- Distributed mode and inline mode of transmissions for CCMs.
- Ethernet link trace (ETH-LT) and loopback (ETH-LB) for CFM sessions.
- Interface Status TLV and remote defect indication (RDI) for the CCM frames on the PS service interface.
- CFM Action profile configuration for events such as adjacency-loss, interface-status-tlv, and RDI on the PS service interface to stop transit traffic.

[See [Ethernet OAM Connectivity Fault Management](#) and [Pseudowire Headend Termination \(PWHT\) Configuration](#).]

- **CFM support on demux interfaces (MX240, MX480, MX960, MX10004, and MX10008)**—Starting in Junos OS Release 24.2R1, we support CFM UP MEP on demux VLAN CCC interface with enhanced CFM mode configuration for both local switching and L2 circuit. This feature enables advanced monitoring and fault detection capabilities for customer VLANs connected through demux VLAN ccc interfaces. It allows for the detection and isolation of faults in the network, ensuring reliable and efficient connectivity.

Additionally, we also support Ethernet link trace and loopback ping on CFM UP MEP on demux interface with enhanced CFM mode. This feature enhances the fault monitoring, fault verification, and fault isolation capabilities of CFM in Ethernet networks.

[See [Configure Connectivity Fault Management \(CFM\)](#) and [Demultiplexing Interface Overview](#).]

- **Support for enhanced request support information (MX204, MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, we've deprecated the CLI option `brief` from the request support information command and introduced the following CLI options to the existing request support information command:

- archive
- with-logs
- with-components
- with-options.

[See [request support information](#).]

- **Support for NETCONF Call Home (MX304, MX960, MX2020, MX10004, MX10008, and MX10016)**—Starting in Junos OS Release 24.2R1, Junos devices support the NETCONF Call Home feature for establishing a NETCONF session over SSH. NETCONF Call Home enables the Junos device to initiate a secure connection to a NETCONF client. You can use NETCONF Call Home when the NETCONF client cannot initiate a connection with the server. This situation can occur when a firewall or another security tool restricts management access to the server or implements Network Address Translation (NAT). NETCONF Call Home can also streamline the initial deployment of network devices by enabling a device to register with a management system when it is first powered on.

[See [NETCONF Call Home](#).]

- **Traffic statistic optimization for PRPD Flex Routes (MX304, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, you can view detailed traffic statistics for PRPD flex routes with the `show programmable-rpd statistics` command.

[See [show programmable-rpd](#) and [show programmable-rpd statistics](#).]

- **Support for IPv4 transit statistics on pseudowire and PWHT interfaces (MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, you can view IPv4 transit statistics for pseudowire and PWHT interfaces using the `show interfaces ps-interface extensive` command.

[See [show interfaces extensive](#).]

- **Support for 64-bit nanosecond EPOCH timestamp over port-mirrored packets (MX304, MX10008, and MX10016)**—Starting in Junos OS Release 24.2R1, you can specify that the software provide a 64-bit nanosecond EPOCH timestamp over a port-mirrored packet for family any packets mirrored in ingress and egress directions.

You set the timestamping feature by using the `packet-timestamp` configuration statement at the `[edit forwarding-options port-mirroring]` hierarchy level.

The port-mirroring destination can be a next-hop group. In this case, every mirrored packet, for each member members of the group, carries the same timestamp.

The timestamp on the mirrored packet is extracted during port-mirror post processing, which executes *after* the mainline packet is processed. Thus there is a microseconds-worth delay since the mainline packet entered or exited on the corresponding interface. Also, an L2 or L3 feature that

depends on the MAC address for forwarding of the mirrored packet might not function as expected, because the MAC header fields are overwritten with the timestamp.

[See [Timestamping of Port-Mirrored Packets](#) .]

- **Support for port mirroring for demux logical interfaces in family CCC (MX240, MX480, MX960, MX10004, and MX10008)**—Starting in Junos OS Release 24.2R1, you can add a demux configuration on top of an ae (aggregated Ethernet) interface configuration with family ccc in global-based or instance-based port mirroring. This feature allows you to use demux interfaces in your port-mirroring configuration to substantially reduce the number of logical interfaces that are consumed by child physical interfaces under the AE bundle.

[See [Applying Layer 2 Port Mirroring to Family ccc Traffic with Demux Logical Interfaces Over Aggregated Ethernet](#) .]

Precision Time Protocol (PTP)

- **Synchronous Ethernet with G.8262 Standard Support (MIC-3D-10GE-SFP-E)**—Starting in Junos OS Release 24.2R1, MIC-3D-10GE-SFP-E for MPC2E-3D and MPC3E-3D line cards of MX240, MX480, MX960, MX2010, and MX2020 routers support synchronous Ethernet features compliant with the following International Telecommunication Union Telecommunication Standardization (ITU-T) standard. This facilitates the transference of clock signals over the Ethernet physical layer:

Synchronous Ethernet (G.8262)—Timing and synchronization aspects in packet networks. Specifies timing characteristics of synchronous Ethernet equipment clock (EEC).

[See [Synchronous Ethernet](#)]

- **Support for Precision Time Protocol with G.8275.1 Standard Support (MIC-3D-10GE-SFP-E)**—Starting in Junos OS Release 24.2R1, MIC-3D-10GE-SFP-E for MPC2E-3D and MPC3E-3D line cards of MX240, MX480, MX960, and MX2020 routers support Precision Time Protocol (PTP) features compliant with the following International Telecommunication Union Telecommunication Standardization (ITU-T) standards. This facilitates the distribution of precise time and frequency over packet-switched Ethernet networks:
 - G.8275.1—PTP profile for phase and time (full timing support)
 - G.8275.1—PTP profile for phase and time over link aggregation group (LAG)

[See [Precision Time Protocol](#)]

- **Support for Precision Time Protocol with Media Access Control Security encryption (MIC-3D-10GE-SFP-E)**—Starting in Junos OS Release 24.2R1, MIC-3D-10GE-SFP-E for MPC2E-3D and MPC3E-3D line cards of MX240, MX480, MX960, and MX2020 routers support Precision Time Protocol (PTP) with Media Access Control Security (MACsec) encryption enabled on the same port at the same time.



NOTE: The following limitations are applicable:

1. The maximum limit for MACsec-enabled logical interfaces (IFL) is 200 per system.
2. The maximum limit for MACsec-enabled ports with physical interfaces (IFDs) and IFLs where MACsec and PTP are enabled together on different ports is 200 per system.
3. The maximum number of IFLs that can be supported on both 1G and 10G ports is 128.
4. PTP in clear text mode is not supported.

[See [Guidelines to Configure PTP over Ethernet](#).]

- **Support for PTP Default profile over IPv4 unicast (MIC-3D-10GE-SFP-E on MX240, MX480, MX960, and MX2020)**—Starting in Junos OS Release 24.2R1, the MX240, MX480, MX960, and MX2020 routers support the PTP default profile (IEEE1588). IEEE1588 defines best master clock algorithm options, configuration management options, and path delay mechanisms (peer delay or delay request-response). The following key features are supported:
 - Support for distribution of phase and time in 1G and 10G speeds.
 - Support for default profile in IPv4 unicast mode.

[See [PTP Profiles](#).]

Public Key Infrastructure (PKI)

- **PKI notifications support for CMPv2 protocol with jsd process (MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, your MX Series router sends public key infrastructure (PKI) notification to Juniper Extension Toolkit (JET) services process (jsd) when it performs certificate management using Certificate Management Protocol (CMPv2) protocol to add, update, and clear certificate operations.

[See [Juniper Extension Toolkit Developer Guide](#).]

Routing Policy and Firewall Filters

- **Support for five tuple match conditions (MX204, MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, MX series routers running Junos OS support five tuple or five combined match conditions for high scale firewall filter performance improvement with only 5-tuple configurations per term without range configurations by avoiding execution of multiple terms in sequence.

[See [fast-lookup-tuple](#), [fast-lookup-tuple-list](#), and [fast-lookup-tuple-list\(policy-options\)](#).]

Routing Protocols

- **Support for OSPFv2 HMAC SHA-2 keychain authentication and weighted ECMP (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-48MP, EX4100-H-12P, EX4100-H-12P-DC, EX4100-H-24P, EX4100-H-24P-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX304, MX150, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, and VMX)**—Starting in Junos OS Release 24.2R1, you can enable OSPFv2 keychain module with HMAC-SHA2 authentication to authenticate packets reaching or originating from an OSPF interface. HMAC SHA2 algorithms include HMAC-SHA2-256, HMAC-SHA2-384 and HMAC-SHA2-512 as defined in RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*. We also support the HMAC-SHA2-224 algorithm. This feature ensures smooth transition from one key to another for OSPFv2 with enhanced security. We also support HMAC-SHA1 and HMAC-SHA2 authentication for virtual and sham links.

You can enable weighted ECMP for directly connected routers. In earlier releases, Junos OS ECMP algorithm does not take the underlying bandwidth into consideration. The algorithm assumes that the links are of equal capacity and the traffic is distributed equally based on this assumption.

To enable OSPFv2 HMAC-SHA2 authentication, configure the keychain *keychain-name* configuration statement [edit protocols ospf area *area-id* interface *interface-name* authentication] at the hierarchy level and algorithm (hmac-sha2-224 | hmac-sha2-256 | hmac-sha2-384 | hmac-sha2-512) option at the [edit security authentication-key-chains key-chain *key-chain-name*] hierarchy level.

To enable keychains authentication support for OSPFv2 virtual links, configure the keychain *keychain-name* configuration statement [edit protocols ospf area *area-id* virtual-link *neighbor-id* router-id transit-area *area-id* authentication] at the hierarchy level.

To enable keychains authentication support for OSPFv2 sham links, configure the keychain *keychain-name* configuration statement [edit protocols ospf area *area-id* virtual-link *neighbor-id* router-id transit-area *area-id* authentication] at the hierarchy level.

To enable weighted ECMP traffic distribution on directly connected OSPFv2 neighbors, configure weighted one-hop statement at the [edit protocols ospf spf-options multipath] hierarchy level.

[See [Understanding OSPFv2 Authentication](#) and [Understanding Weighted ECMP Traffic Distribution on One-Hop OSPF Neighbors](#).]

- **Support for SRLG link constraint in FAD and delay normalization (MX Series)**—Starting in Junos OS Release 24.2R1, we support delay normalization and Flexible Algorithm Definition (FAD) defined constraints related to admin-groups and shared risk link group (SRLG) as defined in RFC 9350, IGP Flexible Algorithm. We also support delay normalization on the listed platforms.

During flexible algorithm computation, when the measured latency values are not equal and the difference is insignificant, IS-IS advertises this slightly higher latency value as a metric. IS-IS uses this normalized latency delay value instead of the measured delay value.

To configure flexible algorithm application specific SRLG values, include the application-specific statement at the [edit protocols isis interface *interface-name* level *level*] hierarchy level.

To exclude SRLG constraint in an FAD, include the exclude-srlg statement at the [edit routing-options flex-algorithm *name* definition] hierarchy level.

[See [delay-measurementlevel](#), and [definition](#).]

- **BGP link bandwidth community (cRPD, EX4100-48MP, EX4300-MP, EX4400-48MP, EX4650, EX9204, EX9208, MX240, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020, cSRX, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200, and QFX5210)**—Starting in Junos OS Release 24.2R1, BGP can communicate link speeds to remote peers, enabling better optimization of traffic distribution for load balancing. A BGP group can send the *link-bandwidth* non-transitive extended community over an EBGp session for originated or received and readvertised link-bandwidth extended communities.

To configure the non-transitive link bandwidth extended community, include the bandwidth-non-transitive: *value* in the export policy at the [edit policy-options community *name* members *community-ids*] hierarchy level.

To enable the device to automatically detect and attach the link-bandwidth community on a route at import, include the [auto-sense](#) auto-sense statement at the [edit protocols bgp group link-bandwidth] hierarchy level. This feature facilitates the integration of devices with different transmission speeds within the network, enabling efficient traffic distribution based on link speed.

[See and [group \(Protocols BGP\)](#).]

- **BMP Improvements (MX10016)**—Starting in Junos OS Release 24.2R1, we have enhanced the robustness and debuggability of BMP to detect root cause problems in BGP connectivity.

To enable collection of time-series data, include in-memory-profiling statement at the [edit routing-options bmp] hierarchy level.

To save advertisement state information, include the keep-advertisement-state statement at the [edit routing-options bmp] hierarchy level.

[See [bmp](#)]

- **Consistent load balancing on flex-algo routes (MX240, MX480, MX960, MX10003, MX10004, MX10008, and MX10016)**—Starting in Junos OS Release 24.2R1, we support consistent hashing on flex-algo routes in a BGP network. You can prevent reordering of flex-algo routes to active paths in an ECMP group when one or more paths fail. BGP overrides the default behaviour of disrupting all existing including active, TCP connections when an active path fails and redirects only inactive flows.

- **Enable RFC 7606 based Error Handling in BGP (MX10016)**—Starting in Junos OS Evolved Release 24.2R1, we support RFC 7606, *Revised Error Handling for BGP UPDATE Messages* that revises the BGP error handling and recommends attributes discard and treat-as-withdraw where the errors can be tolerated instead of a session reset. However, where the errors are too severe, a session reset is triggered. This minimizes the impact of a malformed update message on routing by retaining the established sessions and valid routes.

The `bgp-error-tolerance` statement at the `[edit protocols bgp]` hierarchy level is enabled by default. You can still configure sub-options such as, `malformed-route-limit`, `malformed-update-log-interval`, and `no-malformed-route-limit` under this configuration statement. Note that If you delete the `bgp-error-tolerance` statement, the feature will still remain enabled and the sub-options are reset to their default values.

[See [bgp-error-tolerance \(Protocols BGP\)](#).]

- **FLT on BGP FlowSpec Filters (MX204, MX240, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, you can enable Fast Lookup Table Filter (FLT) to significantly improve packet throughput with BGP FlowSpec. To enable FLT, include the `fast-lookup-filter` statement at the `[routing-options flow]` hierarchy level.

[See [fast-lookup-filter \(Protocols BGP\)](#).]

- **HMAC authentication with hash functions for IS-IS (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-48MP, EX4100-H-12P, EX4100-H-12P-DC, EX4100-H-24P, EX4100-H-24P-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4600-VC, EX4650, EX4650-48Y VC, EX9204, EX9208, EX9214, MX204, MX240, MX304, MX150, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, we extend support to the IS-IS keychain with the following hash functions:
 - HMAC-SHA2-224,
 - HMAC-SHA2-256,
 - HMAC-SHA2-384,
 - HMAC-SHA2-512

Currently, IS-IS supports inline authentication using simple password, keyed MD5 and HMAC-SHA1 algorithms with common keychain. Note that it's important to have the system time synchronized on all nodes when a keychain is active on an IS-IS session.

[See [Understanding Hitless Authentication Key Rollover for IS-IS](#).]

- **Support for BGP VPN to Global RIB Import (cRPD and MX480)**—Starting in Junos OS Release 24.2R1, we support leaking of BGP VPN routes to global RIBs to provide service providers the

flexibility to allow internet access to VPN customers. To configure this feature, include the `vpn-global-import policy` statement at the `[edit routing-options inet.0]` hierarchy level.

To use the auto router discovery feature with router-id without allocating an IP-address include the `route-distinguisher-id-use-router-id` statement at the `[edit routing-options]` hierarchy level.

[See [route-distinguisher-id-use-router-id](#), and [vpn-global-import](#).]

- **Support for configuring multiple independent IGP instances of OSPFv2 (MX204, MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, you can configure and run multiple independent IGP instances of OSPFv2 simultaneously on a router as defined in RFC 6549, *OSPFv2 Multi-Instance Extensions*.

With this feature:

- You can use multiple IGP instances of OSPFv2 to redistribute routes among independent OSPFv2 domains on a single router.
- You can construct flexible OSPFv2 hierarchies across independent IGP domains.
- You can achieve a more scalable OSPFv2 deployment.

To enable multiple IGP instances of OSPFv2 routing on the routing device, configure `ospf-instance igp-instance-name` at the `[edit protocols ospf]`



NOTE: Junos OS does not support configuring the same logical interface in multiple IGP instances of OSPFv2.

[See [Multiple Independent IGP Instances of OSPFv2 Overview](#).]

- **Enhanced IRB Scalability for ARP and ND Unicast Next Hops (MX Series)**—The enhanced IRB scalability feature now supports up to 1.5 million ARP (Address Resolution Protocol) and ND (Neighbor Discovery) unicast next hops per chassis, significantly improving network scalability. To enable this feature, configure the `enhanced-scale` statement at the `[edit interfaces irb unit logical-unit-number]` hierarchy level and the `irb-enhanced-scale` statement at the `[edit bridge-domains bridge-domain-name]` hierarchy level. This configuration allows the Kernel to allocate next hop identifiers from an extended space, ensuring stable and efficient operation. This feature is supported for plain IRB interfaces and does not support advanced functionalities like EVPN-MPLS or VXLAN.

[See [enhanced-Scale \(IRB Interface\)](#) and [irb-enhanced-scale \(Bridge Domain\)](#).]

- **Enhanced MVPN Functionality with Inactive Route Support (MX204, MX240, MX304, MX150, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, and VMX)**—Enhanced MVPN functionality now supports querying inactive routes from shards. This enables MVPN to access and utilize inactive route data for required features. MVPN processes

involving inactive routes are now handled asynchronously, ensuring smoother and more efficient operations.

Securing GTP and SCTP Traffic

- **SCTP DDoS support (MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**

Enhanced support for distributed denial of service (DDoS) filters now extends to the SCTP for advanced forwarding toolkit (AFT) line cards, following its initial deployment for UKERN line cards. Starting from 24.2R1, we segment SCTP packets into two categories: SCTP initialization packets (sctp-init) and unclassified packets (sctp-uncls). DDoS allows direct application of bandwidth, burst, and other filters to SCTP initialization packets. Additionally, users can monitor metrics such as priority, dropped packets, received packets, and violation information for SCTP initialization packets.

DDoS protection filters empower users to handle unexpected surges in traffic directed at the device. Users can define the expected packet bandwidth, priority, and burst rate using DDoS policers. When control traffic exceeds the default or configured policer values, the device drops excess packets and processes the traffic within set limits. Each violation triggers immediate notification, enabling swift response to potential attacks. The device logs each violation, and records the start time and the time of the last observed violation for further analysis.

[See [protocol \(DDoS\)](#)]

Serviceability

- **Automated monitoring of continuous fabric drops (MX240, MX480, MX960, MX2010, and MX2020 routers)**—Starting in Junos OS Release 24.2R1, we support automated monitoring of continuous fabric drops on MX240, MX480, MX960, MX2010, and MX2020 routers. The platform adds a syslog entry with fabric queue information and raises one of these minor errors in the chassis management module:

- CM_CMERROR_FABRIC_LO_PRI_Q_DROP
- CM_CMERROR_FABRIC_LO_PRI_Q_DROP_ALL
- CM_CMERROR_FABRIC_HI_PRI_Q_DROP
- CM_CMERROR_FABRIC_HI_PRI_Q_DROP_ALL

[See [Chassis-Level User Guide](#).]

- **Changes to the show system firmware command to display the versions for the running, primary, and golden firmware for various components on the FRU (MX10004 and MX10008; LC480, LC2101, LC4800, and LC9600 line cards)**—Starting in Junos OS Release 24.2R1, we've enabled new options for the show system firmware command. You can use these options to see firmware version information

for only a particular component rather than all of the components. You can also display detailed firmware version information for all of the components on the field-replaceable unit (FRU).

[See [show system firmware](#).]

- **Support for enhanced request support information (MX204, MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, we've deprecated the CLI option `brief` from the `request support information` command and introduced the following CLI options to the existing `request support information` command:

- `archive`
- `with-logs`
- `with-components`
- `with-options`.

[See [request support information](#).]

Services Applications

- **Inline active flow monitoring multiple BGP next hop support (MX240, MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 24.2R1, we have added support for reporting an accurate BGP next-hop address for the load-balanced traffic over multiple BGP peers in the ingress direction. Prior to this release, we reported the first address in a list of BGP next hops. To contain this accurate BGP next-hop address, we use the IPv4 BGP Nexthop Address (IE 18) field in the IPv4 and MPLS-IPv4 templates and the IPv6 BGP Nexthop Address (IE 63) field in the IPv6 and MPLS-IPv6 templates, for both the IPFIX and the version 9 formats. To configure this feature, include the `multi-bgp-path` statement at the `[edit services flow-monitoring (version-ipfix | version9) template]` hierarchy level. For IPv6 and MPLS-IPv6 flows, you also need to configure the `ipv6-extended-attrib` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level. When this feature is enabled, the `fragmentIdentification` (IE 54) field reports a value of 0.

[See [Understand Inline Active Flow Monitoring](#), [ipv6-extended-attrib](#), and [multi-bgp-path](#).]

- **Inline active flow monitoring support for a demux0 interface mapped to an underlying Aggregated Ethernet (AE) interface, for core facing interfaces only (MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vMX)**—Starting in Junos OS Release 24.2R1, we support configuring inline active flow monitoring for a `demux0` interface mapped to an underlying AE interface, but for only the `mpls-template` of family `mpls` and for protocol family `ccc`. (No `ip4-template` or `ipv6-template` support.) Existing maximum flow scale on the linecard is supported.

Demux0 is a single interface under which you can map an AE interface. In turn, this AE interface hosts multiple child logical interfaces. The flow record reports the SNMP ID of the underlying AE

logical interface as the egress interface. Demux0 logical interfaces are not supported on the reserved unit numbers 16383 and 32767; for example, demux0.16383 and demux0.32767 are not supported.

[See [Understand Inline Active Flow Monitoring](#).]

- **RPM and TWAMP hardware-timestamp and RTT measurement support for tests over a Layer 3 VPN PE-to-PE configuration (MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—Starting in Junos OS Release 24.2R1, you can now configure hardware timestamping on PEs in a Layer 3 VPN PE-to-PE configuration that uses the line cards MPC1 through MPC11E, LC480, LC2101, LC2103, and LC2301. You can also receive round-trip time (RTT) measurements using these line cards across platforms (MPC7 and MPC10).

[See [Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches](#) and [Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#).]

- **Support for Inline IPsec (MX304)**—Starting in Junos OS Release 24.2R1, MX304 supports inline IPsec. The IPsec architecture provides a security suite for the IPv4 and IPv6 network layers. The suite provides such functionality as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. In addition to IPsec, the Junos OS also supports the IKE, which defines mechanisms for key generation and exchange, and manages security associations (SAs).

See [IPsec Overview](#).

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Support for IPv6 endpoints for SR-MPLS DTM SR-TE tunnels (MX10003, MX10004, MX10008, and MX10016)**—Starting in Junos OS Release 24.2R1, we support IPv6 end points for SR-MPLS DTM segment routing-traffic engineering (SR-TE) tunnels. You can configure IPv6 destination networks under spring-te dynamic tunnels and support dynamic segment list and DCSPF (using compute-profile). The following SR-TE dynamic tunnel models are supported:

- IPv6 endpoint for DTM uncolored SR-TE tunnels
- IPv6 endpoint for DTM SR-TE tunnels (SR-MPLS) with inet6color.0 model
- IPv6 endpoint for DTM SR-TE tunnels (SR-MPLS) with transport-rib model

To support transport-rib model for IPv6 DTM SR-TE tunnels, include the `use-transport-class` statement at the `[edit dynamic-tunnels tunnel-name spring-te]` hierarchy level.

If the `use-transport-class` statement is not configured, then, catch all route and application route is created in the `inet6color.0` table. If the `use-transport-class` statement is configured then catch all route and application route is created in `color.inet6.3` table. This behavior is irrespective of using the `use-transport-class` statement at the `[edit protocols source-packet-routing]` hierarchy level. For DT tunnels,

SR-TE takes preference of the `use-transport-class` statement at the `[edit dynamic-tunnels tunnel-name spring-te]` hierarchy rather than at the `[edit protocols source-packet-routing]` hierarchy level.

- **BGP classful transport support for dynamic tunnels and colored transport-rib for next-hop-based tunnels (MX304, MX10004, and MX10008)**—Starting in Junos OS Release 24.2R1, we support colored transport-rib model for next-hop-based dynamic tunnels. By default GRE tunnels are logical interface-based tunnels. IPIP and UDP tunnels are next-hop based tunnels. GRE tunnels can also be configured as next-hop based tunnels by including the `GRE next-hop-based-tunnel` statement at the `[edit routing-options dynamic-tunnels]` hierarchy level.

For logical interface and next-hop-based tunnels, dynamic tunnel specific route addition is triggered when an application route with protocol next-hop is resolved on dynamic tunnel catch-all route.

To support colored transport-rib model for DTM next-hop based tunnels, you should configure the `use-transport-class` statement under the `[edit dynamic-tunnels tunnel-name]` configuration. If the `use-transport-class` statement is not configured then catch all route and application route is created in the `inet(6)color.0` table. If the `use-transport-class` statement is configured then catch all route and application route is created in the `color.inet(6).3` table. If you include the `best-effort` statement at the `[edit routing-options dynamic-tunnels dynamic-tunnel-name destination-networks ip-address]` hierarchy level, dynamic tunnels are created in the `inet(6)color.0` table.

To enable the `use-transport-class` statement under dynamic tunnel, you should include `auto-create` statement at the `[edit routing-options transport-class]` hierarchy level.

To configure colored transport-rib, you should include the `preserve-nexthop-hierarchy` statement at the `[edit routing-options resolution]` hierarchy level.

Software Installation and Upgrade

- **Base OS update (ACX710, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, Junos OS uses the FreeBSD main base OS. This upgrade provides improved security and better performance. In earlier releases, Junos OS used the FreeBSD Release 12 base OS.

[See [Junos® OS Software Installation and Upgrade Guide](#).]

- **In-band ZTP management in campus fabrics (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX9204, EX9208, EX9214, MX304, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can simplify the provisioning process for campus fabrics. Already provisioned upstream devices, such as core and distribution devices, that are capable of detecting downstream Day 0 devices can provide Layer 3 connectivity. With Layer 3 connectivity, the downstream Day 0 devices can proceed with Secure ZTP.

To configure in-band ZTP management, enable the `in-band-ztp` statement at the `[edit system services]` hierarchy on your core and distribution devices. Optionally, your cloud controller can provide the in-band-ztp configuration as part of the provisioning process for your core and distribution devices.

See [Zero Touch Provisioning](#)

- **Migration of Linux OS version**—Starting in Junos OS Release 24.2R1, the following devices support Wind River Linux LTS22:

Table 6: List of devices that support Wind River Linux LTS22

Platforms	Routing Engines Supported
ACX5448, ACX5448-D, and ACX5448-M	RE-ACX-5448
EX9204, EX9208, and EX9214	EX9200-RE2
MX240, MX480, and MX960	RE-S-X6
MX2010, MX2020	REMX2K-X8
MX2008	REMX2008-X8-64G
MX10008, MX10004	JNP10K-RE1
MX204	MX204
MX10003	JNP10003-RE1
MX304	JNP304-RE
SRX1600	SRX1600
SRX2300	SRX2300
SRX5800, SRX5600, and SRX5400	SRX5K-RE3

Table 6: List of devices that support Wind River Linux LTS22 (Continued)

Platforms	Routing Engines Supported
QFX10002-60C	RE-QFX10002-60C

Starting in Junos OS Release 24.2R1, in order to install VM Host image based on Linux WR LTS22, you have to upgrade the i40e NVM firmware version to 9.1 or later.

Subscriber Management and Services

- **L2TP Tunnel Switch Accounting Support for Network Gateway on MPC10E (MX240, MX480, MX960), MX10K-LC9600 (MX10004 and MX10008), and MX304**—Starting in Junos OS Release 24.2R1, the following line cards support accounting features of L2TP Tunnel Switched (LTS) sessions:
 - MPC10E for MX240, MX480, MX960.
 - MX304 with integrated line card.
 - MX10K-LC9600 for MX10004 and MX10008 routers.

The LTS accounting supports monitoring that includes:

- Interim and final statistics reported to RADIUS server.
- CLI based statistics queries.

The outputs of the following commands display LTS statistics related to the line cards:

- `show subscriber accounting-statistics id sid`
- `show services l2tp tunnel-switch tunnel statistics`

See [[show subscribers](#) and [show services l2tp tunnel-switch tunnel](#).]

- **Support for CCC family on demux interfaces (MX240, MX480, MX960, MX10004, and MX10008)**—Starting in Junos OS Release 24.2R1, circuit cross-connect (CCC) family is supported on demux interfaces. This allows you to use demux interfaces for switching and layer 2 circuit configurations.

See [[Configuring CCC Encapsulation for Layer 2 VPNs](#), [IP Demultiplexing Interfaces](#), and [Subscriber Interfaces and Demultiplexing Overview](#).]

Additional Features

We have extended support for the following features to these platforms.

- **Symmetric IRB with EVPN Type 2 routes** (MX Series). You can enable EVPN Type 2 routing over symmetric integrated routing and bridging (IRB) interfaces in an EVPN-MPLS edge-routed bridging (ERB) overlay.

[See [Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics](#).]

- **Support for ISSU with slice-based HCoS** (MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)

[See [Hierarchical Class of Service for Network Slicing](#).]

- **Firewall filter support on demux interfaces** (MX240, MX480, MX960, MX10004, and MX10008)

- Interface-specific filters attached on inet, inet6, ccc, and any firewall filter families
- Firewall filter attachment on a demux interface in input and output directions
- Policers attachment on demux interface in input and output directions
- Flexible firewall filter match conditions
- Two-color and three-color policers

[See [Subscriber Interfaces and Demultiplexing Overview](#) and [Firewall Filters Overview](#).]

- **G.8275.2 (PTS) profile BC and OC support with LAG (IPv4 and IPv6) on MX10008 Universal Routing Platform** (MX Series) [G.8275.2 Enhanced Profile](#)

What's Changed

IN THIS SECTION

- [EVPN | 96](#)
- [Flow-based and Packet-based Processing | 97](#)
- [General Routing | 97](#)
- [Junos OS API and Scripting | 99](#)
- [Platform and Infrastructure | 100](#)
- [Subscriber Access Management | 100](#)
- [User Interface and Configuration | 100](#)
- [VPNs | 101](#)

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

EVPN

- **OISM SBD bit in EVPN Type 3 route multicast flags extended community**—In EVPN Type 3 Inclusive Multicast Ethernet Tag (IMET) route advertisements for interfaces associated with the supplemental bridge domain (SBD) in an EVPN optimized intersubnet multicast (OISM) network, we now set the SBD bit in the multicast flags extended community. We set this bit for interoperability with other vendors, and to comply with the IETF draft standard for OISM, draft-ietf-bess-evpn-irb-mcast. You can see this setting in the output from the `show route table bgp.evpn.0 ? extensive` command.

[See [CLI Commands to Verify the OISM Configuration](#).]

- **Group-based Policy (GBP) tag displayed with `show bridge mac-table` command**—On platforms that support VXLAN-GBP, the `show bridge mac-table` command now displays a GBP TAG output column that lists the GBP tag associated with the MAC address for a bridge domain or VLAN in a routing instance. Even if the device doesn't support or isn't using GBP itself, the output includes this information for GBP tags in packets received from remote EVPN-VXLAN peers.

[See [Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN](#).]

- **Default behavior changes and new options for the easy EVPN LAG configuration (EZ-LAG) feature**—The easy EVPN LAG configuration feature now uses some new default or derived values, as follows:
 - Peer PE device `peer-id` value can only be 1 or 2.
 - You are required to configure the loopback subnet addresses for each peer PE device using the new `loopback peer1-subnet` and `loopback peer2-subnet` options at the `[edit services evpn device-attribute]` hierarchy level. The commit script uses these values for each peer PE device's loopback subnet instead of deriving those values on each PE device. These replace the `loopback-subnet` option at the `[edit services evpn device-attribute]` hierarchy level, which has been deprecated.
 - If you configure the `no-policy-and-routing-options-config` option, you must configure a policy statement called `EXPORT-LOO` that the default underlay configuration requires, or configure the new `no-underlay-config` option and include your own underlay configuration.
 - The commit script generates "notice" messages instead of "error" messages for configuration errors so you can better handle `[edit services evpn]` configuration issues.
 - The commit script includes the element names you configure (such as IRB instance names and server names) in description statements in the generated configuration.

This feature now includes a few new options so you have more flexibility to customize the generated configuration:

- `no-underlay-config` at the `[edit services evpn]` hierarchy level—To provide your own underlay peering configuration.
- `mtu overlay-mtu` and `mtu underlay-mtu` options at the `[edit services evpn global-parameters]` hierarchy level—To change the default assigned MTU size for underlay or overlay packets.

[See [Easy EVPN LAG Configuration](#).]

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the `[edit protocols evpn]` hierarchy level. In most use cases, you don't need to change the default limit. If you want to change the default limit, we recommend that you don't set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

[See [mac-ip-limit](#).]

Flow-based and Packet-based Processing

- The subscription path for the flow sensor is changed from `/junos/security/spu/flow/usage` to `/junos/security/spu/flow/statistics`. This change maintains a uniform path in request and response data.

General Routing

- **Change in the XML tags displayed for the `show virtual-network-functions` command in JDM (Junos node slicing)** — To align the XML tags displayed for the `show virtual-network-functions gnf-name | display xml` with the new XML validation logic, we have replaced the underscores (`_`) in the output with hyphens (`-`) as shown below:

Old output:

```
user@host> show virtual-network-functions mgb-gnf-d | display xml (vnf-instance) 1mgb-gnf-dRunningdown <ip_addr>192.168.2.1</ip_addr> <<< The tag includes _ 2 <max_mem>16GiB</max_mem> <<< The tag includes _ <resource_template>2core-16g</resource_template> <<< The tag includes _ <qemu_process_id>614702</qemu_process_id> <<< The tag includes _ <smbios_version>v2</smbios_version> <<< The tag includes _
New output: user@jdm> show virtual-network-functions mgb-gnf-d | display xml 1mgb-gnf-dRunningdown192.168.2.1 <<< The tag changes to ip-addr. 216GiB <<< The tag changes to max-mem. 2core-16g <<< The tag changes to resource-template. 614702 <<< The tag changes to qemu-process-id. v2 <<< The tag changes to
```

smbios-version. This change is applicable to any RPC that previously had underscores in the XML tag name.

- When you run the `run show lldp local-information interface interface-name | display xml` command, the output is displayed under the `lldp-local-info` root tag and in the `lldp-local-interface-info` container tag. When you run the `run show lldp local-information interface | display xml` command, the `lldp-tlv-filter` and `lldp-tlv-select` information are displayed under the `lldp-local-interface-info` container tag in the output.
- Change in use of RSA signatures with SHA-1 hash algorithm? Starting in Junos OS Release 24.2R1, there is a behavioural change by OpenSSH 8.8/8.8p1. OpenSSH 8.8/8.8p1 disables the use of RSA signatures with SHA-1 hash algorithm by default. You can use RSA signatures with SHA-256 or SHA-512 hash algorithm.
- Starting from Junos OS 21.4R1 platforms with the following Routing Engines which have Intel CPUs with microcode version 0x35 observe the error warning, "000: **Firmware Bug:** TSC_DEADLINE disabled due to Errata; please update microcode to version: 0x3a (or later)" on the console. RE-S-X6-64G RE-S-X6-128G REMX2K-X8-64G RE-PTX-X8-64G RE-MX2008-X8-64G RE-MX2008-X8-128G
- **Non-revertive switchover for sender based MoFRR**— In earlier Junos releases, source-based MoFRR ensured that the traffic reverted to the primary path from the backup path, when the primary path or session was restored. This reversion could result in traffic loss. Starting in Junos OS 22.4R3-S1, source-based MoFRR will not revert to the primary path, i.e. traffic will continue to flow through the backup path as long as the traffic flow rate on the backup path does not go below the configured threshold set under protocols `mvpn hot-root-standby min-rate`.

[See [min-rate](#)]

- **Show active forwarding session for sender based MoFRR**— The `show multicast route extensive` command will show the active forwarding session in the case of source-based MoFRR. The field `Session Status: Up and Forwarding` will indicate that the particular session is currently forwarding traffic.

[See [show multicast route](#).]

- For MPC5E line card with flexible-queuing-mode enabled, queue resources are shared between scheduler block 0 and 1. Resource monitor CLI output displays an equal distribution of the total available and used queues between scheduler blocks. This correctly represents the queue availability to the routing engine.
- Change to the commit process—In prior Junos OS and Junos OS Evolved releases, if you use the `commit prepare` command and modify the configuration before activating the configuration using the `commit activate` command, the prepared commit cache becomes invalid due to the interim configuration change. As a result, you cannot perform a regular commit operation using the `commit` command. The CLI shows an error message: 'error: Commit activation is pending, either activate or

clear commit prepare'. If you now try running the commit activate command, the CLI shows an error message: 'error: Prepared commit cache invalid, failed to activate'. You then must clear the prepared configuration using the clear system commit prepared command before performing a regular commit operation. From this Junos and Junos OS Evolved release, when you modify a device configuration after 'commit prepare' and then issue a 'commit', the OS detects that the prepared cache is invalid and automatically clears the prepared cache before proceeding with regular 'commit' operation.

[See [Commit Preparation and Activation Overview](#).]

- In a firewall filter configured with a port-mirror-instance or port-mirror action, if l2-mirror action is also configured, then port-mirroring instance family should be any. In the absence of the l2-mirror action, port-mirroring instance family should be the firewall filter family.
- **G.8275.1 profile configuration with PTP, SyncE, and hybrid mode (Junos)**— On all Junos platforms, when configuring the G.8275.1 profile, it is mandatory to configure Precision Time Protocol (PTP), Synchronous Ethernet (SyncE), and hybrid mode. Earlier, the system would not raise a commit error even if the required hybrid and SyncE configurations were missing while configuring G.8275.1 profile. However, going forward you will not be able to configure the G.8275.1 profile without configuring PTP, SyncE and hybrid mode to be compliant with the ITU-T standards.

[See [G.8275.1 Telecom Profile](#).]

Junos OS API and Scripting

- **Changes to the XML output for ping RPCs (MX480)**—We've updated the junos-rpc-ping YANG module and the corresponding Junos XML RPCs to ensure that the RPC XML output conforms to the YANG schema. As a result, we changed the XML output for the following ping RPCs:
 - <ping>—The XML output emits <ping-error-message> and <ping-warning-message> tags instead of <xnm:error> and <xnm:warning> tags.
 - <request-ping-ce-ip>—The XML output is enclosed in an <lsping-results> root element.
 - <request-ping-ethernet>—
 - The <ethping-results> root tag includes a <cfm-loopback-reply-entry> or <cfm-loopback-reply-entry-rapid> tag for each received response. In earlier releases, a single tag enclosed all responses.
 - The XML output includes only application specific error tags and omits <xnm:error> tags.
 - The <cfm-loopback-reply-entry-rapid> tag is now reflected in the YANG schema.
 - <request-ping-overlay>—The <ping-overlay-results> element includes a new child tag <hash-udp-src-port>.

Platform and Infrastructure

- Starting Junos Evolved Release 24.1R1, support for Network Time Protocol (NTP) over TLS (RFC 8915 compliant) for the ACX-series, QFX-series, and PTX-series includes:
 - Support to configure local-certificate for server and certificate verification option for client.
 - Verification of x.509 certificates to establish a TLS channel between client and server.
 - TLS NTS-KE protocol support
 - Support for NTS secured client-server NTP communication at server and client.
 - Support for new NTS options in commands "system ntp nts", "system ntp server (server_name) nts remote-identity", and "show ntp associations no-resolve" commands.

Subscriber Access Management

- You can configure VLAN termination cause codes to specify RADIUS attribute values for different termination scenarios on JUNOS MX Series platforms supporting the Layer-2 Bitstream Access (L2BSA) feature. You can diagnose and manage network issues effectively by understanding the specific reasons for VLAN termination. Ensure that the correct termination cause codes are sent by validating configuration and testing scenarios to correctly interpret network events. When a subscriber logs out, the system occasionally sends an incorrect termination cause value to RADIUS. The subscriber VLAN "Account-Terminate-Cause" in "Acct-Stop" message for different L2BSA subscriber logout error scenarios is modified to display correct reasons for termination.

[See "VLAN Termination Causes and Code Values" and "show network-access aaa terminate-code".]

User Interface and Configuration

- Viewing files with the `file compare files` command requires users to have maintenance permission**—The `file compare files` command in Junos OS and Junos OS Evolved requires a user to have a login class with maintenance permission.

[See [Login Classes Overview](#).

- Configuration database maximum size increased (ACX Series, EX Series, MX Series, QFX Series, SRX Series, and vSRX)**—We've enhanced the `extend-size` statement at the `[edit system configuration-database]` hierarchy level to increase the maximum database size. On devices with a default configuration database size of ~400 MB, `extend-size` increases the maximum database size to ~2 GB. On devices with a default configuration database size of ~660 MB, `extend-size` increases the maximum database size to ~2.2 GB.

[See [configuration-database](#).]

VPNs

- **Increase in revert-delay timer range**— The revert-delay timer range is increased to 600 seconds from 20 seconds.

[See [min-rate](#).]

- **Configure min-rate for IPMSI traffic explicitly**— In a source-based MoFRR scenario, you can set a min-rate threshold for IPMSI traffic explicitly by configuring `ipmsi-min-rate` under `set routing-instances protocols mvpn hot-root-standby min-rate`. If not configured, the existing min-rate will be applicable to both IPMSI and SPMSI traffic.

[See [min-rate](#).]

Known Limitations

IN THIS SECTION

- General Routing | [101](#)
- Layer 2 Ethernet Services | [102](#)
- Platform and Infrastructure | [103](#)
- Services Applications | [103](#)

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

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

General Routing

- When do a `snmpwalk` on chassisd-related mib (for example: `jnxOperatingTable`) on MX104, the chassisd cpu may go up to 70%. In MX104 device total number of entries fetched during this walk is about 1000 entries and above. This `snmpwalk` takes more time to complete the SNMP polling. Due to MX104 available memory and processor we could see chassisd spike during `snmpwalk`.

[PR1604901](#)

- The time delay on GNF reconnect is due to socket drop on alarmd and chassis went for a reconnect, during this reconnect process the ready messages from GNFs are ignored until the chassis reconnect timer is expired. The socket drop on alarmd and reconnect is expected while performing GRES. [PR1771319](#)
- The error message Received IFD attach for interface name with PIC in Offline or Offline wait state.Ignore it , IFFPC: ifd detach returned error 7 is usually observed when a PIC reconfiguration is triggered while the IFDs are still getting installed at the linecard. PIC reconfiguration is mainly due to a PIC mode change. Before triggering a pic mode change, in addition to checking the state of the pic mode i.e if the PIC is online, also verify that the IFDs are attached correctly either by allowing adequate delay between PIC bounce (approximately 3 to 5 minutes) or observing logs.[PR1774974](#)
- The error message is observed when a PIC reconfiguration is triggered while the IFDs are still getting installed at the linecard. PIC reconfiguration is mainly due to a pic mode change. Before triggering a pic mode change, in addition to checking the state of the pic mode i.e if the PIC is online, also verify that the IFDs are attached correctly either by allowing adequate delay between PIC bounce (~3-5mins) or observing logs.[PR1780251](#)
- If there is i2cs upgrade failure , do not run online reload command. please re run the jfirmware upgrade. in case if you executed online reload command after I2cs upgrade failure please jack out jack in the line card before running the jfirmware upgrade again. [PR1783364](#)
- If SGRP over subscription is configured for a BNG-UP port which hardware doesn't support over subscription it still accepts the port into the SGRP. The result is subscribers are not handled properly. [PR1791676](#)
- 1PPS performance drop is seen during clksyncd process restart with 1588 default profile. Its a baseline Mx Aloha line card behaviour. [PR1796244](#)
- show system license output does not display hardware based licensing details. show system license bandwidth flex-only output displays linecard specific bandwidth details. It makes the calculation very complicated to add advance/premium details to this output. Tier can be chassis, linecard or port based. Bandwidth details is for all ports of the linecard irrespective of the tier of each port.[PR1797309](#)
- 40g interface doesnt support EM policy feature but it will still display in the cli output of show chassis temp-threshold as it gets created as "et" interface. [PR1807219](#)

Layer 2 Ethernet Services

- Issue was seen when test was done back to back GRES within 5mins time. this is expected behavior from the system as per current architecture. Suggestion is to wait for sometime before maybe 10 minutes or so for subsequent GRES. [PR1801234](#)

Platform and Infrastructure

- When global level changes are made like changing Mac-age, with large scale and high CPU a watchdog core is generated without any functional impact. The config change is pushed from l2alm to pfe, which loops through all the RTTs and associated bd and ifbd. If the process is not yielded within 240sec, the pfeman watchdog is called in and writes a core. [PR1775966](#)

Services Applications

- In Junos OS Release 17.4 and later, subscriber sessions on the LNS that send an ICRQ that includes RFC5515 AVPs may fail to establish a session. The client will receive a CDN error "receive-icrq-avp-missing-random-vector" in response. [PR1493289](#)

Open Issues

IN THIS SECTION

- [General Routing | 104](#)
- [Layer 2 Ethernet Services | 108](#)
- [Network Management and Monitoring | 109](#)
- [Platform and Infrastructure | 109](#)
- [Routing Protocols | 109](#)
- [Services Applications | 110](#)
- [User Interface and Configuration | 111](#)

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

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

General Routing

- You might see a misleading syslog message "L2CKT/L2VPN acquiring mastership for primary" although no VPN/L2CKT is configured on the router. [PR1105459](#)
- For the MPC10E card line, the IS-IS and micro-BFD sessions do not come up during baseline. [PR1474146](#)
- The Sync-E to PTP transient simulated by Calnex Paragon Test equipment is not real network scenario. In real network deployment model typically there will be two Sync-E sources (Primary and Secondary) and switchover happens from one source to another source. MPCE7 would pass real network SyncE switchover and associated transient mask. [PR](#)
- When the active slave interface is deactivated, the PTP lock status is set to 'INITIALIZING' state in show ptp lock-status output for few seconds before BMCA chooses the next best slave interface. This is the day-1 behavior and there is no functional impact. [PR1585529](#)
- Percentage physical-interface policer is not working on aggregated Ethernet, after switching between baseline configuration to policer configuration. [PR1621998](#)
- On all Junos OS platforms, agentd process crash will be seen in telemetry streaming longevity test. [PR1647568](#)
- There will be drop of syslog packets seen for RT_FLOW: RT_FLOW_SESSION_CREATE_USF logs until this is fixed. This will not impact the functionality. [PR1678453](#)
- %PFE-x: fpcx user.err ppmman: [Error] PPM:PROCESSOR_L2TP_SF: PpmProcProtoL2tpSf::processPkt: No tunnel entry found for received L2TP tunnel control packet. LocalAddr: x.x.x.x LocalTunnelId: 0 Timestamp xx:xx:xx device fpcX user.err ppmman: [Error] PPM:PROCESSOR_L2TP_SF: PpmProcProtoL2tpSf::processPkt: Received packet ipv4 header parsing failed. PacketSize:xx [PR1689921](#)
- On Junos OS platforms, even though there are no active subscribers, a foreign file propagation (ffp) commit error is seen for the class-of-service traffic-control-profile. [PR1700993](#)
- When LAG is configured with mixed speed interfaces switching to a secondary interface of different port speed, results in a few packet drops for a very short duration. PTP remains lock and there is no further functional impact. [PR1707944](#)
- Once the device is loaded with the new image, PIC tries to boot up. mspmand is one of the processes inside PIC, crashes sometimes. [PR1714416](#)
- fec-codeword-rate data with render type decimal64 is rendered as string in grpc python decoder. [PR1717520](#)

- JDI-RCT:M/Mx: SMPC crash @ hostif_clear_toe_interrupts, toe_interrupt_handler after fpc restart scenario .[PR1733053](#)
- PTP state went to Freerun and acquiring before phase-aligning again when the SyncE ESMC is disabled or downgraded from GM or the upstream node one hop above the parent node.
[PR1738532](#)
- On MXVC , Due to some timing issue when RPD is restarted, It will not be spawned again. This issue is rarely reproducible.[PR1740083](#)
- On all Junos OS and Junos OS Evolved platforms BGP traceoptions configuration will have an impact on the CPU, threads will be busy and will take time to recede in spite of disabling it. It is important we enable a specific trace flag and disable it when the CPU goes high. It is also important not to perform switchover and other triggers which can add load to the CPU during traces are enabled. Traces must be enabled discretely.[PR1724986](#)
- Error message might occur once in a while with full scale during negative scenarios like 'clear bgp neighbor all' with all the services like EVPN, vrf etc being present.[PR1744815](#)
- RBU DIAGS REGRESSIONS: MX480 CommonDiag::JDE3(volt_services_show_clients) failing on MPC7e. [PR1747033](#)
- RBU DIAGS REGRESSIONS: MX2010 Diagnostics::Jde3Diag(phy_reg_access) test is failing.
[PR1747297](#)
- On all MX Series platforms, faulty hardware issue on MIC due to clock sync error generated brings down the interfaces without any major alarm or log notification.[PR1749943](#)
- MX480: Observed multiple na-mqttd.core-tarball@mosquitto_send_suback,mqtt3_handle_subscribe,mqtt3_packet_handle. [PR1758264](#)
- On Junos MX Series platforms with Trusted Platform Module (TPM), reset of master password got stuck post device reboot.[PR1760822](#)
- On MX Series platform with a combination of MPC1-9, LC480, LC2101, and MPC10E, MPC11E, LC9600 line cards, when preserve-nexthop-hierarchy knob enabled and maximum-ecmp configured with more than 32 next-hops in the MPLS FRR (Multiprotocol Label Switching fast-Reroute) and BGP (Border Gateway Protocol) Multipath scenario, packet loss when primary path is added back in ECMP nexthop (say after primary interface or session is marked UP) will be higher compared to that on MX platform with MPC1-9, LC480, LC2101 line cards only, OR with MPC10E, MPC11E, LC9600 line cards only. This packet loss is proportional to the value in maximum-ecmp configuration.[PR1765856](#)
- On Junos OS platforms, when executed just after line-cards are up after system-reboot, the CLI output for active-errors (show system errors active, show system errors active detail) displays empty output for some initial duration that can run for minutes. Issue is seen when number of errors present in line-card is very high (10K+). Since all these errors need registration with CLI serving

daemon running on Routing Engine, before it can display error-info, CLI output for this command is delayed. However, as a workaround alternative CLI (`show chassis errors active detail`) can be used, which displays similar output.[PR1775073](#)

- MX10008 :: PLD is higher than 2000 msec on ungraceful removal of a fabric board.[PR1776054](#)
- Commit error is needed when streaming server and export-profile is not configured properly. With the incomplete configuration that is missing below might cause the interfaces to go down upon reboot of the unit or FPC. `set services analytics streaming-server profile name remote-address ip set services analytics streaming-server profile name remote-port port set services analytics export-profile profile-name reporting-rate rate`. This needs to be greater than 1.[PR1779722](#)
- Even after `request vmhost power-off` LEDs keep lighting on. The LEDs state should be off because routing-engine doesn't have power in case of `request vmhost power-off`.[PR1781815](#)
- The `show` command cause performance degradation or hog CPU. [PR1784219](#)
- V6 Endpoint SRTE: 4PE (IPv4 over IPv6) routes in inet.0 table are not getting resolved in inet6color table because 4PE is not supported with inet(6)color model. 4PE can be supported with transport class. [PR1786029](#)
- When interfaces with different speed are configured as members of AE, some of the members are not added to aggregated Ethernet. And if GRES is enabled, vmcore might be generated on backup Routing Engine. [PR1799451](#)
- Under scaled configurations with interfaces, FW filters, routing protocols etc certain script based config/unconfig automated operations, in the presence of continuous traffic, can encounter one or more PPE traps that momentarily cause traffic drops. These momentary traffic drops happen at the tail end of the time the business configuration is removed and a baseline configuration is loaded in a single commit. These do not cause any service or functionality impact. [PR1800967](#)
- On MX Series platforms with SCBE3-MX (MX240, MX480 and MX960) due to a hardware failure of the Control Board, the Routing Engine(RE) switchover might not happen. This will result in the 19.4Mhz clock failure and has potential risk for chassis wide traffic impact. In worst case all revenue ports will be impacted. If the RE switchover is done in a timely manner then the device will recover because FPCs will try using the 19.4Mhz clock from the new primary.[PR1801284](#)
- MPC11 ISSU command fails from Junos OS Release 24.1R1 to 24.2R1 and causes MPC11 linux crash. The issue only applies to ULC image.[PR1803205](#)
- This issue is caused because of the fact that peers-synchronize is configured, and master-password is configured to encrypt the config being sync'ed. However, as there is no master-password configured on the peer device, the encrypted configuration cannot be decrypted (this is expected). This has not been supported from day-1, however a workaround can be done in order to get this to work. The workaround is to manually configure the same master password on the peer device manually. At a high level the problem is as follows: Consider there are two devices A and B in a peer-sync config 1.

config on dev A contains secrets which need to be encrypted with the master password and synced with the device B 2. The master-password (juniper123+masterpassword) is configured on device A and the configuration is encrypted and written to /tmp/sync-peers.conf 3. The /tmp/sync-peers.conf is then synced to device B but device B does not have the same master-password configured which results in the config failing to decrypt. The master-password itself is not a part of the config-database. Additionally, it cannot be transmitted over an unencrypted HA Link, as this would lead to the master-password getting leaked. This is by design, and would be a security concern if it were to be transmitted across an unencrypted channel. Therefore, this work as designed. In order to work around this issue follow these steps: 1. configure the master-password on device B and commit the config 2. configure the same master-password on device A and commit the config and it should get sync'ed correctly.[PR1805835](#)

- It's day-1 non-functional issue, where in PFE the unilist is coming with proto as IPv4 for IPv6 unilist nexthop. [PR1806717](#)
- PLL Access Failure alarms is observed on a MPC11E line card of REV 53 after loading 24.2I-20240429.0.0958 on a MX2010 box[PR1808044](#)
- LLDP neighbor does not recover after protocol is enabled globally on the router.[PR1811545](#)
- Under scaled configurations with interfaces, FW filters, routing protocols etc, certain script based config/unconfig automated operations in the presence of continuous traffic, can encounter one or more PFE traps that momentarily cause traffic drops. These momentary traffic drops happen at the tail end of the time the business configuration is removed and a baseline configuration is loaded in a single commit. These do not cause any service or functionality impact. [PR1800967](#)
- When LC4800 is operated in the worst case operating corner, i.e. all ports running at full line rate, NEBS ambient temperature = 55C, high altitude, there is a possibility that the PCIe switch temp sensor on the SIB8 (JNP10008-SF2) can falsely report a yellow alarm for over temperature. This issue is applicable to Junos 24.2R1 and 24.2R1-S1 releases. Hardware Symptoms tracking signature: cli show system alarms Alarm time Class Description 2024-03-07 02:28:45 PST Minor SFB 2 PCIe Switch Temp Sensor Warm. [PR1801778](#)
- On MX platforms with LC2101 line cards and 10-gigabit ethernet interfaces configured in loopback mode, when Line card is booted multiple times, the ethernet interfaces on line card remains down and traffic on those interfaces will be impacted. [PR1809511](#)
- With 24.2R1 software release, some of the 100G and 400G links might remain DOWN after LC4800 FPC restart. Check workaround for recovery. [PR1814101](#)
- On all Junos OS platforms with Border Gateway Protocol (BGP) rib-sharding enabled and NSR(Nonstop Active Routing) configured, upon deactivation and activation of routing-instances, interfaces and protocols together, memory leak in rpd is observed. There will be no traffic impact at this time because memory leak is very slow. [PR1761191](#)

- We've noticed that we are experiencing 100% CPU utilization during GRES in the smid and alarmd daemons. After investigating, we discovered that the gRPC client (feature daemon) was sending RPC requests and waiting for responses while the gRPC server (license-check daemon) was down during GRES. If the client doesn't receive a response within the specified deadline time (30 seconds), it should return a GRPC_QUEUE_TIMEOUT type and success value as 0. But, it got stuck in getting "grpc_core::Timestamp::Now ()" API. [PR1805723](#)
- On MX platforms when LC4800, the fan speed at 25C ambient temperature may exceed 44% of the max speed target. Therefore the system may not meet NEBS acoustic requirements. [PR1824343](#)

Layer 2 Ethernet Services

- In order to allow protocol daemons (such as rpd, dot1xd et. al.) to come up fast when master password w/ TPM is configured, the daemons must be allowed to cache the master-password when they read their configuration. In order to cache the master-password, the daemons must individually reach out to the TPM to decrypt the master password and cache it in their memory. This scenario leads the TPM to be flooded with decryption requests, and therefore causes the TPM to be busy and start rejecting decryption requests. To prevent the daemons from core dumping in this scenario, and to allow successful decryption of secrets, we retry the decryption request to the TPM. However, to allow the TPM queue to drain, we introduce a sched_yield() call before retrying to sleep for 1 quantum of time. Without this, we will fail on all our retries. Additionally, a decryption request can also take a large amount of time (> 5 secs). This results in SCHED_SLIP messages being seen in the logs, as the requesting process is idle while the decryption request is being processed by the TPM. This can exceed the SCHED_SLIP timeout, and result in libjtask logging the SCHED_SLIP messages into the configured system log file. These SCHED_SLIPs should not cause any route instability, are benign, and can be ignored as these are seen only during configuration consumption by the various daemons. [PR1768316](#)
- DHCP asymmetric-lease-time is slow processing large scale requests to terminate 64K subscribers. This condition applies to both DHCP local server and DHCP relay when asymmetric-lease-time is configured. Regardless of the timer value configured the JDHCPD process will only handle 20 request per second which results in longer than expected time to terminate all DHCP subscribers. In DHCP dual stacked environments the client termination is split between protocol type, 10 clients for DHCPv4 and 10 clients for DHCPv6 are terminated per second. In the example of having 64K dual stacked subscribers with minimum asymmetric-lease-time of 600, after network disruption, there is a 600 second interval for detection JDHCPD takes an approximate addition 53 minutes to terminate all 64K subscribers. The engineering fix for this PR will process 100 client request per second rather than the original 20 requests per second. [PR1817227](#)

Network Management and Monitoring

- In some NAPT44 and NAT64 scenarios, Duplicate SESSION_CLOSE Syslog will be seen. [PR1614358](#)

Platform and Infrastructure

- PCT-VIRTUAL: Firewall filter counters are not incremented as expected when filter is applied to IRB interface in the ingress/egress direction via forwarding table [PR1766471](#)
- As per the current cos design, we aren't merging the configuration from wildcard and explicit configurations, instead explicit configs takes precedence and we don't apply the wildcard configs. For example: set class-of-service interfaces xe-* scheduler-map sch0 ---> explicit xe-* takes precedence set class-of-service interfaces all unit 0 classifiers dscp cls set class-of-service interfaces xe-* unit 0 classifier dscp cls set class-of-service interfaces xe-3/2/0 unit 0 classifier dscp cls2 -> explicit interface config takes precedence set class-of-service interfaces xe-3/2/0 unit * classifier dscp cls set class-of-service interfaces xe-3/2/0 unit 0 rewrite_rule rw_rule -> takes precedence, classifier won't be applied It's recommended to add the required config also along with the explicit config, like in this reported case, we need to below config to fix this problem. set class-of-service interfaces xe-* scheduler-map sch0 set class-of-service interfaces xe-* unit 0 classifiers dscp cls [PR1797119](#)
- Few Error message may occurs while deleting multiple EVPN ETREE Routing Instances. [PR1808643](#)

Routing Protocols

- LDP OSPF are in synchronization state because the IGP interface is down with ldp-synchronization enabled for OSPF. user@host> show ospf interface ae100.0 extensive Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity. As per the current analysis, the IGP interface goes down because although LDP notified OSPF that LDP synchronization was achieved, OSPF is not able to take note of the LDP synchronization notification, because the OSPF neighbor is not up yet. [PR1256434](#)
- On MX Series platforms, unexpected log message will appear if the CLI command 'show version detail' or 'request support information' is executed: test@test> show version detail *** messages *** Oct 12 12:11:48.406 re0 mcsnoopd: INFO: krt mode is 1 Oct 12 12:11:48.406 re0 mcsnoopd: JUNOS SYNC private vectors set. [PR1315429](#)

- On Junos platforms and Junos Evolved platforms, if a BGP peer goes down and stays down, the system might take an extremely long time to complete removing the BGP routes. The issue is observed when a BGP peer sends many routes, only a small number of routes are selected as the active routes in the routing information base (RIB, also known as the routing table), and if the BGP delete job gets only a small part of the CPU time because other work in the routing process is utilizing the CPU.[PR1695062](#)
- On all Junos OS platforms and Junos Evolved with scaled BFD sessions, FPC reload/restart results in few BFD session flap.[PR1698373](#)
- Memory leaks in rt_instance block rpd process when deactivating and then activating protocols, routing-instances and interfaces. [PR1761191](#)
- With BGP sharding and NSR configured, deactivating or activating routing-instances and interfaces back to back multiple times on active Routing Engine leads to generate rpd core files on backup Routing Engine at rt_flash_queue_insert. [PR1781293](#)
- Configuration of a global AS number is necessary when route target filter is enabled. Currently JUNOS cli does not enforce configuring a global AS number and it has been the behavior for a long time. Many unexpected issues may be seen without a global AS number. It's been a recommended practice to configure a global AS number in the field.[PR1783375](#)
- On all Junos OS platforms having BGP (Border Gateway protocol) configured, when route is leaked through rib-group from one routing instance to another having the same AS (Autonomous System) number and one of the routing-instances has BGP configured with local-as, it is observed that even after configuring loops with any value greater than one as the number of loops option, the route still remains hidden instead of being active which results in traffic drop. [PR1771344](#)

Services Applications

- On Junos MX80, MX240, MX480, MX960 platforms with Multiservices Modular Interfaces Card (MS-MIC), Multiservices Modular Port Concentrators (MS-MPC) service cards, in an issue where an old dynamic security association_configuration (sa_cfg) for a tunnel is present and trying to establish new sets of Internet Protocol Security Security Association (IPSec SAs) using a new Internet Key Exchange (IKE) SA established for the same remote device but with a different request. This can happen, if for some reason old sa_cfg is not cleaned (failed in clean-up). On crash, the Key Management Daemon (kmd) restarts but fails because of kernel instance mismatch present in the kernel database. So all the IPsec tunnels will be impacted.[PR1771009](#)

User Interface and Configuration

- To recover from this and to avoid problem due to problem delta synchronize, "set system commit no-delta-synchronize" can be configured as work-around (no-delta-synchronize is hidden knob but safe to use). It will enforce entire ?juniper.conf? to synchronize rather than delta changes and will help in this case. [PR1801136](#)
- Upgrade from Junos OS Release 20.3x to 24.2R1 fails if extend-db config stanza is present This issue is happening due to extend-db knob configured in the config. Delete the extend-db knob, reboot the box and then perform the upgrade. Issue is not seen. [PR1806109](#) and [PR1807931](#)

Resolved Issues

IN THIS SECTION

- [Application Layer Gateways \(ALGs\) | 112](#)
- [Class of Service \(CoS\) | 112](#)
- [EVPN | 112](#)
- [General Routing | 113](#)
- [Interfaces and Chassis | 118](#)
- [Junos Fusion Provider Edge | 118](#)
- [Layer 2 Ethernet Services | 118](#)
- [MPLS | 119](#)
- [Network Management and Monitoring | 119](#)
- [Platform and Infrastructure | 119](#)
- [Routing Policy and Firewall Filters | 120](#)
- [Routing Protocols | 120](#)
- [Subscriber Access Management | 120](#)
- [User Interface and Configuration | 121](#)
- [VPNs | 121](#)

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

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

Application Layer Gateways (ALGs)

- The SIP UDP register packet will be dropped on MX platforms. [PR1781379](#)

Class of Service (CoS)

- An error was observed while performing the interface flap by making PIC on/off. [PR1793588](#)
- COS output-traffic-control-profile is not getting attached to ps transport logical interfaces on a specific MX Series platforms that support HCOS. [PR1795898](#)
- Traffic impact during ISSU across FPCs on Junos MX platforms [PR1796770](#)
- xSTP does not work in ephemeral-db mode on all Junos OS Evolved platforms after the l2cpd restart [PR1800645](#)
- On Junos OS platforms the telemetry subscribe to path: "/components/component[name='Routing Engine0']/state/memory/used" is not working as expected. [PR1800754](#)
- Filter will be configured with incorrect vlan-IDs and commit error will not be displayed. [PR1802362](#)
- XMCHIP PFEs could run into XMCHIP_CMERROR_CPQ_INT_REG_QSYS_QUEUE_UNDRN_ERROR during ungraceful SIB or Peer-FPC power off event or due to bad fabric links. [PR1807812](#)
- [MX] daemon.err rshd[618008]: Second port outside reserved range. [PR1807939](#)
- L2ALD core is observed due to presence stale IFD. [PR1810013](#)

EVPN

- DF election hold timer stopped if there is a remote T4 route received already. [PR1787293](#)
- Deleting/adding EVPN Instance and interface having ESI configuration resulting in traffic loss. [PR1789065](#)
- Traffic loss is seen in multihoming EVPN-MPLS single-active scenario. [PR1790705](#)

- Traffic drop in between EVPN hosts might be seen when `vlan-id none` is configured and any commit changes done under `[edit protocols mpls]` hierarchy. [PR1792566](#)
- MHPE shows as `I_ESI` instead of `RNVE` in EVPN-VXLAN configurations without the DCI interconnect configuration. [PR1799761](#)
- The VXLAN traffic drop could be seen after modifying control-word in an EVPN instance [PR1807084](#)

General Routing

- LTS19: MX960: 000: [Firmware Bug]: TSC_DEADLINE disabled due to Errata; please update microcode to version: 0x3a (or later) seen upon upgrade to Junos OS Release 21.4. [PR1608045](#)
- LR4 and SR4 100G optics, every JPN should have a generic description. [PR1718902](#)
- CRPD log file is growing indefinitely and not being rotated by JCNr. [PR1727111](#)
- The pkid process failed during restart. [PR1729592](#)
- PLL core frozen alarm might be seen in rare occasion. [PR1742266](#)
- MX : Input Feed state to be displayed in SNMP `jnxOperatingState`. [PR1742996](#)
- MX10004/MX10008 - DIP Switch - 15A or 20A - Not displayed in the CLI output and has to be checked physically. [PR1744396](#)
- Telemetry data is not exported in an IS-IS scaled Segment Routing scenario. [PR1745615](#)
- The chassisd logs indicate incorrect logs when fan trays are removed on MX10000 platforms. [PR1753787](#)
- FTF egress filter is not supported. [PR1756572](#)
- The license-check can get restarted. [PR1760259](#)
- Multiple Routing Engine switchover can cause IDEEPROM failure on FPC and PSM. [PR1760978](#)
- With BGP sharding, observed memory leak in cookie `ifx_dist_msg`. [PR1761238](#)
- During ISSU , observed errors in MPC10 card @ WANIO core: WI backpressure caused packets to get errored or discarded for 100G core 4. [PR1765931](#)
- The FPC Crash will be observed on MX Series platforms. [PR1766578](#)
- The SFB3 will go offline during SFB3-ADC-MPC7E link initialization. [PR1768592](#)

- At MX2020 with SFB3 and MPC7+ADC in FPC slot 11, link errors in fabric planes 1, 4, 7, 10 will be seen upon MPC7/ADC restart in FPC slot 11. [PR1769983](#)
- xSTP configured interface will remain in discarding state. [PR1770053](#)
- XML VALIDATION Failed for cli: clear interfaces transport pm. [PR1771024](#)
- DUT is sending same source-port-id for two PTP primary links connected to downstream node with multiline card scenarios. [PR1772138](#)
- Connectivity between static and LDP signaled pseudowires broken in VPLS after upgrade and results in VPLS traffic drop. [PR1772424](#)
- **/components/component/power-supply/** sensor leaves and values are not streamed in gNMI and the data is streamed through gRPC in latest primary Junos OS Release 24.2. [PR1772435](#)
- The demux0 interfaces won't show any data traffic (measured in packets per second) after the FPC restarts. [PR1773148](#)
- Behaviour of unsupported subscription modes for IP AFT prefix filtering feature. [PR1773412](#)
- Observed aft-ulcd core on MPC11E line card while doing ISSU on Junos OS 23.1 and later releases. [PR1773571](#)
- One TEMP sensor is updated through the display snmp mib walk jnxFruTemp function. All four and five temperature sensors (ABPM) must be updated. [PR1775383](#)
- Internet traffic destined for the Network Address Translation (NAT) pool address will loop after configuration changes. [PR1776355](#)
- Interface stay in link DOWN state when using third party optics. [PR1776596](#)
- Host device cannot resolve target IP's ARP when client uses virtual MAC address. [PR1776782](#)
- BGP session on MPLSoUDP tunnel will not be established. [PR1776783](#)
- The evo-aftmand-bt process crash is observed during an Routing Engine switchover. [PR1776828](#)
- ARP resolution fails when EVPN configured with preserve-nexthop-hierarchy. [PR1776913](#)
- PIM-RP registration fails in PIM EVPN gateway configuration scenario. [PR1777493](#)
- On MX Series platform replacing the line-cards may trigger FPC to be offlined due to unreachable destinations. [PR1777534](#)
- RPD core can be seen when running Telemetry for protocols from top of tree and doing routing instance add/ delete operation. [PR1778103](#)
- The FPC connection times out and reboots post mastership switchover. [PR1778324](#)

- After a GRES the MX Series external clock is stuck in a holdover state. [PR1781161](#)
- Junos OS Release 23.4R1:SFW:commit should fail after deleting existing service-set ss1 and adding a new service set ss2 on a same interface with SFW configuration, but commit success which is unexpected behavior. [PR1781264](#)
- Software upgrade with force option tries to recover space for /var partition and the upgrade fails. [PR1781632](#)
- After Routing Engine switchover PPPoE subscribers may fail to login. [PR1782239](#)
- Delays can be seen while the upgrading process runs due to the status of UFS set to mode enable. [PR1783119](#)
- License key is not installed after USB upgrade, via set system license keys key.. [PR1783509](#)
- The mspmand process might crash on the MS-MPC during deletion of service-sets configuration. [PR1783745](#)
- RPC error is observed when deleted operation fails while deleting configuration through GNMI-client. [PR1783817](#)
- L2VPN: On AFT-based cards, Using Core facing IRB, l2vpn datapath broken. [PR1783821](#)
- FPC reboot seen on MX Series platforms with PMB memory correctable errors. [PR1784080](#)
- Devices with MXVC+SPC3 service card experience failure in NAT pool allocation when configuration for balancing network traffic (AMS LB) is used. [PR1784696](#)
- On Junos MX Series platforms silent FPC reboot is observed with no generation of crash files. [PR1785182](#)
- On MX Series platform replacing MPC2E-NG with MPC3E-NG will trigger destination errors and cause other MPC7/MPC10 FPCs went offline. [PR1785241](#)
- The rpd crash is observed due to segment fault. [PR1785884](#)
- MPC line card crashes while ISSU to Junos OS Release 24.1 or later, displays "ISSU PREPARE TIMEOUT" error. [PR1785960](#)
- Seeing rep-serverd core while running the testcase1 for script dynamic_ae_vlandemux_dualStack_dhcp_relay_mts_mx.robot on CUPS-L2 regression. [PR1786253](#)
- TARGET DEFINED subscriptions to some paths will not send periodic data. [PR1786663](#)
- Autoneg configuration for 1G Optical is not turning on upon inspecting PHY registers. [PR1787154](#)
- Interface configured with BPDU-disable goes down during VC mastership switchover. [PR1787892](#)

- CFM sessions configured on It interface hosted on a few line cards will not be up. [PR1788491](#)
- LC9600: Continuous Fabric Request timeout errors maybe logged in /var/log/messages when a previously active PFE is disabled due to a fatal fault. [PR1788846](#)
- Watchdog SPI transaction is causing the interface flap in the system. [PR1789272](#)
- Error "Cannot set interface down action on more than one session running on interface et-x/x/x.x rmep x". [PR1790156](#)
- Traffic drop will be seen when two firewalls actions are referring to same Routing Instance name. [PR1790331](#)
- EBGp sessions established over IPsec tunnels would flap if multihop configuration statement with ttl=1 is configured. [PR1791196](#)
- SSH public key of host is base64 encoded twice in bootstrap complete message. [PR1791905](#)
- Flabel sequencing is not correct leading to traffic loops and congestion in the Packet Forwarding Engine on MX Series platforms. [PR1792173](#)
- With any configuration change or interface up or down with MACsec protocol configured with or without Dot1x, dot1xd process core file is observed in the device. [PR1792507](#)
- License is lost on NG-RE platforms after device/routing-engine reboots. [PR1792672](#)
- ASIC-level memory leak occurs and error messages are displayed when when changing or deleting fast-lookup-filters (FLT filters) interface-specific filter. [PR1793044](#)
- Multicast traffic black-holing upon MoFRR primary link went down. [PR1793196](#)
- IRB configuration is not mandatory in EZ-LAG. [PR1793346](#)
- Physical interfaces are not created when back to back port configuration is committed without much delay. [PR1793763](#)
- The 'no-flow-control' interface configuration not retained post ISSU on MX Series platforms. [PR1793999](#)
- Application restart after switchover. [PR1794769](#)
- Port goes down after adding interface configuration and changing the port from 1g copper to 10g fiber. [PR1794939](#)
- JUNOS_REG: MX : With the GR interface configured, ASIC error at PFE can trigger vmcore on backup. [PR1795218](#)
- The broadband subscriber (L2BSA subscribers) on the core interface logging out with interface state changes. [PR1796125](#)

- High CPU utilization due to BMP to be running with the longest and highest run count. [PR1796530](#)
- Traffic impact during ISSU across FPCs on Junos OS MX Series platforms. [PR1796770](#)
- ICMP Echo reply dropped by Junos MX Series platforms with nat port block-allocation. [PR1796974](#)
- We might observe repd core (in the "from" release) during ISSU. There are no functional impact due to this repd core. [PR1797189](#)
- DCPFE process crash occurs in EVPN-VXLAN scenario. [PR1797516](#)
- More than 8 forwarding-classes cannot be configured with CBF. [PR1797592](#)
- BGP learning or convergence performance degradation with BGP RIB Sharding. [PR1797996](#)
- Linecard with [inline-services flex-flow-sizing configuration statement enabled and large scaled routes. [PR1798466](#)
- Routing Engine switchover will cause multiple issues on MX304. [PR1798511](#)
- MPLS/RSVP LSP self-ping behaviour differs from expected behaviour causing self-ping time out. [PR1798801](#)
- JDI-REGRESSION:MX304: cty core seen during usb upgrade from 23.1R1.8 to 23.2R2.18 [PR1799431](#)
- Clksyncd core file is observed when performing snmp mib walk. [PR1800134](#)
- On Junos platforms the telemetry subscribe to path": "/components/component[name='Routing Engine0']/state/memory/utlized is not working as expected. [PR1800754](#)
- IKE is not coming up with dhgroup19 and dhgroup20. [PR1801201](#)
- CPU usage gets spiked for eventd due to flooding of pfe_khms_spurious_wakeup log. [PR1801535](#)
- Filter will be configured with incorrect vlan-IDs and commit error will not be displayed. [PR1802341](#)
- Disable RSA signatures using the SHA-1 hash algorithm by default in Evo OpenSSH module. [PR1802441](#)
- Commit changes do not go through when logical interfaces is set and COS is enabled. [PR1805716](#)
- Telemetry statistics for af interfaces are not getting streamed for UKERN based Linecard. [PR1805769](#)
- Traffic flooding occurs when deactivating and activating interfaces in EVPN scenario [PR1803898](#)
- Observed "MCNHMBB index availability" alarm when system is subjected to heavy PIM join/prune churn. [PR1792740](#)

- The KRT queue stuck resulting in subscriber traffic loss. [PR1797305](#)
- XMCHIP PFEs could run into XMCHIP_CMERROR_CPQ_INT_REG_QSYS_QUEUE_UNDRN_ERROR during ungraceful SIB or Peer-FPC power off event or due to bad fabric links. [PR1807812](#)
- CPU utilization of the rpd process stays high on all Junos OS platforms. [PR1808463](#)
- Traffic loss occurs if persistent link error is seen on a fabric plane to PFE after restarting or rebooting another FPC in a different slot. [PR1808923](#)

Interfaces and Chassis

- The chassis-control subsystem will not respond when fxp0 interface has static ARP configured. [PR1777137](#)
- Scaled interface configurations do not get deleted through openconfig delete. [PR1785035](#)
- The ifinfo process crash is seen on Junos OS platforms. [PR1786555](#)
- CHASSISD_IFDEV_RTSLIB_FAILURE: ifdev_create: rtlib_ifdm_add failed (No such file or directory) after creating a virtual interface tunnel. [PR1798681](#)

Junos Fusion Provider Edge

- The sdp process crashes when trying to add a new satellite device to the network. [PR1787147](#)

Layer 2 Ethernet Services

- The jdhcpd process crash will be observed due to double free of memory allocation when DHCP ALQ is configured. [PR1769598](#)
- All AEs bundles configured in Active-Standby mode for EVPN-MPLS routing-instances will flap on the first commit post a fresh system reboot. [PR1783793](#)
- DHCPv6 Information-request packets will be dropped when it has server identifier option. [PR1779273](#)
- The client session is logging out as DHCP renewal is not successful. [PR1801142](#)

MPLS

- RSVP incorrectly determines the outgoing interface resulting in the rpd crash. [PR1785214](#)
- To add support for egress address tlv for nil fec validation. [PR1795558](#)
- FPC restart can be seen on Junos MX platforms in a telemetry scenario with SRTE configuration. [PR1789901](#)
- RPD process crashes due to memory exhaustion. [PR1793982](#)
- Label Switched Path (LSP) traffic drop observed on a transit router with graceful-restart configured after a brief outgoing link flap. [PR1800034](#)
- An invalid remote neighbor address is shown for IPv6 LDP with session-protection configured. [PR1801384](#)

Network Management and Monitoring

- YANG: After upgradation s/w version on DUT, yang package with lower revisions are available in upgraded software version. [PR1693646](#)
- Ifmd fails to send notification about CRC error is seen on link. [PR1769373](#)
- File descriptor leak in eventd after crash of the standby Routing Engine. [PR1783320](#)
- Warning messages seen while Custom Yang Package Deletion. [PR1804856](#)

Platform and Infrastructure

- Traffic drop is observed with "preserve-nexthop-hierarchy" at global level on Junos MX platforms with MPC10 and above line cards (including MPC10, MPC11, LC9600 line cards and MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and MX304 platforms). [PR1766080](#)
- The vlan tag in arp request packet is not getting removed with vlan-id none configured in EVPN scenario. [PR1794023](#)
- DNS and NTP may not be working as expected on junos 23.3 version above. [PR1795068](#)
- Wedge condition is seen on MX platforms with MPC10/MPC11/JNP10K-LC9600 and MX304 platform switch port utilization is above 80%. [PR1800623](#)

Routing Policy and Firewall Filters

- The rpd crash while adding "from family" for srte template. [PR1798465](#)

Routing Protocols

- Routing loop will be observed during ISIS-FRC implementation. [PR1760334](#)
- The rpd process crashes after multiple iterations of disable/enable IS-IS protocol. [PR1777702](#)
- Multicast traffic loss observed when IGMP/MLD snooping is toggled on a VPLS routing-instance. [PR1781059](#)
- The rpd process keeps on restarting on all Junos and Junos OS Evolved platforms due to memory exhaustion. [PR1781138](#)
- PE Routers continue to drop prefixes in a L3 VPN scenario. [PR1785231](#)
- When the interface is down the MTU for OSPF3 will be the calculated value from interface MTU. [PR1787982](#)
- The Valid routes for BGP are showing invalid with RPKI Origin Validation. [PR1792703](#)
- The rpd crash is observed when PIM SSM mode with RPF-Vector and MoFRR is configured. [PR1792886](#)
- LDP label advertisement is stopped for approximately 25 seconds when micro loop avoidance is enabled in SR over OSPF. [PR1793148](#)
- Multicast Traffic loss (20s - 30s) is seen during the reconvergence after link flap. [PR1793598](#)
- [BGP-RPKI] | no reason marked for prefixes with AS_SET in the path that are rendered invalid. [PR1797395](#)
- [BGP][ipv6] BGP multipath selects wrong interface with "Multiple Single-Hop EBGP Sessions on Different Links Using the Same IPv6 Link-Local Address". [PR1807504](#)

Subscriber Access Management

- Address preservation for delegated prefixes doesn't work for subscribers in VRF. [PR1777967](#)

User Interface and Configuration

- SSH configuration changes do not come into affect on an existing outbound ssh client connection. [PR1791814](#)
- The mgd core seen on MX960 during upgrade to aggregated satellite switches. [PR1799277](#)

VPNs

- The MVPN traffic starts dropping after Routing Engine switchover. [PR1747703](#)
- Higher revert-time for MVPN Hot-root-standby. [PR1777865](#)
- TEF-FT: helper information can be added for newly added ipmsi-min-rate configuration statement. [PR1788769](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 24.2R1 | 122](#)
- [Procedure to Upgrade Junos OS | 122](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 124](#)
- [Upgrading a Router with Redundant Routing Engines | 125](#)
- [Downgrading from Release 24.2R1 | 126](#)

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 24.2R1



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 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-
x86-32-20.4R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-
x86-64-20.4R1.9-signed.tgz
```

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

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-
x86-32-20.4R1.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-
x86-64-20.4R1.9-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 24.2R1, 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:

- 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 sixty months after the first general availability date and customer support for an additional six more months.



NOTE: The sixty months of support for EEOL releases is introduced in Junos OS 23.2 release and is available for all later releases. For releases prior to 23.2, the support for EEOL releases continues to be thirty six 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 7: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	60 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 24.2R1

To downgrade from Release 24.2R1 to another supported release, follow the procedure for upgrading, but replace the 24.2R1 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 | 127](#)
- [What's Changed | 128](#)
- [Known Limitations | 129](#)
- [Open Issues | 129](#)
- [Resolved Issues | 130](#)
- [Migration, Upgrade, and Downgrade Instructions | 131](#)

What's New

IN THIS SECTION

- [Dynamic Host Configuration Protocol | 127](#)
- [Platform and Infrastructure | 127](#)
- [Software Installation and Upgrade | 128](#)
- [VPNs | 128](#)

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

Dynamic Host Configuration Protocol

Platform and Infrastructure

- **Support for faster file transfer (NFX150, NFX250, and NFX350)**—Starting in Junos OS Release 24.2R1, NFX series devices support faster file transfer or file copy. NFX series devices support the following types of file copy:
 - Remote file copy—Use request `vmhost remote-file-copy` command to copy a file directly between the hypervisor and an external fileserver.
 - Local file copy of a file on the hypervisor—Use request `vmhost local-file-copy` command to copy a file or a directory present under `/var/public/` directory structure on the hypervisor to a different name under `/var/public/` directory structure.

See [Configuring VNFs on NFX350 Devices](#) and [Configuring VNFs on NFX250 NextGen Devices.](#)

- **Support for Virtual Port Peering (NFX 250 and NFX350)**—

Starting in Junos OS Release 24.2R1, you can set up mapping or peering between the L2 interface (xe-0/0/x or ge-0/0/x port) and the L3 interface (ge-1/0/0 to ge-1/0/9) by configuring the `virtualization-options interfaces L3-interface-name mapping peer-interfaces physical-interface` at the `[edit vmhost]` hierarchy.

See [Configure Interfaces and VLANs for a VNF](#) .

- **Support to configure static MAC address for Virtual Port Peering (NFX350)**—Starting in Junos OS Release 24.2R1, you can configure static MAC address on the VNF interface while using the Virtual

Port Peering (VPP) feature by configuring interfaces `interface-name mac-address mac-address` at the `[edit virtual-network-functions VNF-name]` hierarchy.

See [Configure Interfaces and VLANs for a VNF](#).

- **Support for new Junos OS user permission (NFX150, NFX250, NFX350)—**

Starting in Junos OS Release 24.2R1, a new Junos user permission, `vnf-operation` makes the request `virtual-network-functions` CLI hierarchy available to Junos OS users that do not belong to root or to super-user class.

You can add the user permission to a custom user class using the statement `vnf-operation` at `[edit system login class custom-user permissions]`.

See [request virtual-network-functions](#).

Software Installation and Upgrade

- **Base OS update (ACX710, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, Junos OS uses the FreeBSD main base OS. This upgrade provides improved security and better performance. In earlier releases, Junos OS used the FreeBSD Release 12 base OS.

[See [Junos® OS Software Installation and Upgrade Guide](#).]

- **Support for USB Autoinstallation (NFX150, NFX250, and NFX350)**—Starting in Junos OS Release 24.2R1, you can autoinstall the Junos OS image by inserting a USB flash drive containing the image and configuration into the USB port of NFX series devices. You can autoinstall the image by configuring the `set system services usb-auto-install` command.

[See [Junos® OS Software Installation and Upgrade Guide](#).]

VPNs

- **Support for Internet Key Exchange Protocol daemon (iked) (NFX Series Routers)**—Starting in Junos OS Release 24.2R1, NFX150 and NFX250 devices support the `iked` process. The `iked` is an Internet Key Exchange (IKEv2) daemon, which performs mutual authentication, and establishes and maintains IPsec flows and security associations (SAs) between the two peers.

See [IP Security on NFX Devices](#).

What's Changed

Learn about what changed in this release for NFX Series devices.

Known Limitations

There are no known limitations in hardware or software in this release for NFX Series devices.

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

Open Issues

IN THIS SECTION

- [General Routing](#) | 129
- [High Availability \(HA\) and Resiliency](#) | 130
- [Interfaces](#) | 130

Learn about open issues in this release 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 NFX platforms, when one partition supports a Junos OS Release 23.4R1 image (supported on LTS19 operating system) and the other partition supports an image older than Junos OS Release 23.4R1 (supported on WRL8 operating system), the request `vmhost reboot disk` command is not executed as expected.

As a workaround, upgrade both the partitions with same image versions [PR1753117](#).

- On the NFX350 devices, `srxpfe` core is seen. [PR1792616](#).

High Availability (HA) and Resiliency

- When high availability (HA) is enabled and fabric links are configured on NFX devices (NFX150, NFX250 and NFX350 with nfx-3 software package), the fabric link monitored status is displayed as Down leading to an FL status.[PR1794559](#)

Interfaces

- On the NFX350 device, even though the ethernet cable is physically plugged in and the `show interface` command displays Front panel LED status as up, the front panel LED is not ON.[PR1702799](#)

Resolved Issues

IN THIS SECTION

- [Interfaces | 130](#)
- [Network Address Translation | 131](#)
- [VNF | 131](#)

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

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

Interfaces

- Starting in Junos OS Release 24.2R1 release, when you run the command `show chassis alarm` on NFX 350 devices, the output displays Major TSensor 3:Coretemp Access Failed due to swapping of the symlinks of `hwmon0` and `hwmon1`.[PR1769699](#)

Network Address Translation

- On the NFX devices when the NAT port number or the IP address of the peer device located behind a Network address translation (NAT) device is changed, the next Dead Peer Detection (DPD) or rekey process fails to update the port number in the existing tunnel NAT Traversal (NAT-T) flow session. The failure to update the port-number happens if the DPD is configured as `always-send`. This condition leads to communication failure over the Internet Protocol Security (IPsec) tunnel.

[PR1776216](#)

VNF

- On the NFX platforms, the `pfe` (Packet Forwarding Engine) process crashes when configured with custom mode templates like flex mode or any other custom mode due to memory exhaustion.

[PR1776815](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 132](#)
- [Basic Procedure for Upgrading to Release 24.2 | 133](#)

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



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

Upgrade and Downgrade Support Policy for Junos OS Releases

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 sixty months after the first general availability date and customer support for an additional six more months.



NOTE: The sixty months of support for EEOL releases is introduced in Junos OS 23.2 release and is available for all later releases. For releases prior to 23.2, the support for EEOL releases continues to be thirty six 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 8: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	60 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 24.2

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 24.2R1:

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 QFX Series

IN THIS SECTION

- [What's New | 134](#)
- [What's Changed | 144](#)
- [Known Limitations | 148](#)
- [Open Issues | 148](#)
- [Resolved Issues | 149](#)
- [Migration, Upgrade, and Downgrade Instructions | 151](#)

What's New

IN THIS SECTION

- [Chassis | 135](#)
- [EVPN | 135](#)
- [Junos OS API and Scripting | 139](#)
- [Junos Telemetry Interface | 139](#)
- [MPLS | 140](#)
- [Network Management and Monitoring | 141](#)
- [Routing Protocols | 141](#)
- [Software Installation and Upgrade | 142](#)
- [Additional Features | 143](#)

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

To view features supported on the QFX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Release 24.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- [QFX10002](#)
- [QFX10008](#)
- [QFX10016](#)
- [QFX10002-60C](#)

Chassis

- **Support for thermal health check and PSM watchdog features (QFX10002)**—Starting in Junos OS Release 24.2R1, thermal health check is supported on QFX10002 switches. You can configure an action to be taken on detection of a thermal health event such as power leakage. You can enable PSM watchdog feature as well.
[See [thermal health check](#) and [watchdog \(PSM\)](#).]

EVPN

- **Default discard policy for GBP filters (EX4100, EX4400, EX4650, and QFX5120)**—Starting in Junos OS Release 24.2R1, you can configure group-based policy (GBP) firewall filters with a default discard policy that is applicable when a packet fails to meet any of the match conditions.
[See [Example: Micro and Macro Segmentation Using Group Based Policy in a VXLAN](#).]
- **MAC/IP inter-tagging for GBP filters (EX4100, EX4400, EX4650, and QFX5120)**—Starting in Junos OS Release 24.2R1, you can apply media access control (MAC)-based GBP firewall filters to routed traffic and IP-based GBP firewall filters to switched traffic. This is called inter-tagging. Previously, MAC-based GBP filters applied to switched traffic and IP-based GBP filters applied to routed traffic. By enabling inter-tagging, your MAC-based and IP-based GBP filters apply to both switched and routed traffic.
[See [Example: Micro and Macro Segmentation Using Group Based Policy in a VXLAN](#).]
- **GBP tag propagation using EVPN Type 5 route advertisements (EX4400, EX4650, and QFX5120)**—Starting in Junos OS Release 24.2R1, we support group-based policy (GBP) tag propagation using EVPN Type 5 route advertisements of IP prefixes. Switches and routers typically use EVPN Type 5 advertisements for exchanging routes between data centers. Prior to this release, we supported EVPN Type 2 to Type 5 route conversion between data centers, which resulted in /32 IP routes being exchanged instead of IP prefix routes.
[See [Example: Micro and Macro Segmentation Using Group Based Policy in a VXLAN](#).]
- **Access security support in EVPN-VXLAN overlay networks (EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, we support access security features on certain EX Series and QFX Series switches that function as Layer 2 VXLAN gateways in an Ethernet VPN–Virtual Extensible LAN (EVPN-

VXLAN) centrally-routed overlay network (two-layer IP fabric). We support the following features on Layer 2 server-facing interfaces that are associated with VXLAN-mapped VLANs:

- DHCPv4 and DHCPv6 snooping [See [DHCP Snooping](#).]
- Dynamic ARP inspection (DAI) [See [Understanding and Using Dynamic ARP Inspection \(DAI\)](#).]
- Neighbor discovery inspection (NDI) [See [IPv6 Neighbor Discovery Inspection](#).]
- IPv4 and IPv6 source guard [See [Understanding IP Source Guard for Port Security on Switches](#).]
- Router advertisement (RA) guard [See [Understanding IPv6 Router Advertisement Guard](#).]

The access security features function the same and you configure them in the same way in an EVPN-VXLAN environment as you do in a non-EVPN-VXLAN environment. However, keep these differences in mind:

- We do not support these features on multihomed servers.
- These features do not influence the VXLAN tunneling and encapsulation process.
- **L2PT with Q-in-Q over VXLAN tunnels in EVPN-VXLAN bridged overlay networks (EX4100-48P, EX4400-48F, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, we support Layer 2 protocol tunneling (L2PT) with Q-in-Q for traffic from an access interface to VXLAN tunnel destinations in a bridged overlay (BO) Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) network. You can use this feature with service-provider style or enterprise-style access interface configurations.

You can use L2PT over VXLAN tunnels with all of the Q-in-Q use cases we support for an EVPN-VXLAN network. For Q-in-Q, the device tunnels tagged frames over VXLAN using the VNI of the VLAN in the frame, and tunnels untagged frames using the VNI of the native VLAN.

To enable this feature, configure the `l2pt` statement at the `[edit protocols layer2-control]` hierarchy level with the access interface name `interface name` and the following required options:

- `destination vxlan-tunnel`—Enable L2PT for traffic toward a VXLAN tunnel destination.
- `protocol protocol-name`—Specify a protocol to tunnel. Include additional `protocol` statements for each protocol you want to tunnel.

[See [Layer 2 Protocol Tunneling over VXLAN Tunnels in EVPN-VXLAN Bridged Overlay Networks, Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network](#), and [l2pt \(Destination Tunnels\)](#).]

- **Suppress EVPN Type 5 host routes from DCI to DC (EX4400-24MP, EX4400-48F, MX304, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can suppress EVPN Type 5 host route

advertisements that re-originate from the data center interconnect (DCI) to the local DC. You can achieve better scaling and performance on leaf devices with this feature.

[See [suppress-host-routes-from-dci-to-dc](#).]

- **Enhanced OISM in EVPN-VXLAN ERB overlay networks with an IPv6 underlay (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can configure enhanced optimized intersubnet multicast (OISM) for IPv4 and IPv6 multicast data traffic with an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) edge-routed bridging (ERB) overlay network that has an IPv6 underlay. To configure this feature:
 - Set up the EVPN-VXLAN fabric with an IPv6 underlay:
 - You can use either external BGP (EBGP) or OSPFv3 with IPv6 addressing for the IPv6 underlay.
 - Use the `inet6` option when you set the VXLAN tunnel endpoint (VTEP) source interface to the device loopback interface in the EVPN instance (EVI):

```
set routing-instances evpn-instance-name vtep-source-interface lo0.0 inet6
```

- Configure the enhanced OISM elements for your multicast EVPN-VXLAN environment in the same way you would configure these elements in an EVPN-VXLAN network with an IPv4 underlay.

You can configure any of the supported platforms as enhanced OISM server leaf devices, and only EX4650 and QFX5120 switches as enhanced OISM border leaf devices.

[See [EVPN-VXLAN with an IPv6 Underlay](#) and [Optimized Intersubnet Multicast in EVPN Networks](#).]

- **Statically identify multihoming peer OISM leaf devices in an EVPN-VXLAN network running enhanced OISM (EX4100-24T, EX4300-MP, EX4400-24MP, EX4400-24P, EX4400-48F, EX4400-48MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can statically configure a multihoming peer leaf device to identify its peers in an Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) network running enhanced optimized intersubnet multicast (OISM). EVPN-VXLAN provider edge (PE) leaf devices are multihoming peers when they share an Ethernet segment (ES) for a multihomed host. With enhanced OISM, if the leaf devices have static information about their multihoming peers, they can avoid multicast traffic loss when their peer devices go down and up again.

On each multihoming peer leaf device, to identify the device's multihoming peers, configure the `multihoming-peer-gateways [peer-device-IPv4-address ...]` statement at the `[edit protocols evpn]` hierarchy level. Specify a list of peer addresses within square brackets, or specify a single peer address without any brackets.

[See [Statically Identify Multihoming Peers With Enhanced OISM To Improve Convergence.](#)]

- **EVPN-VXLAN DCI multicast support with enhanced OISM (EX4400-24MP, EX4400-24P, EX4400-48F, EX4400-48MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, we support Ethernet VPN–Virtual Extensible LAN (EVPN–VXLAN) to EVPN–VXLAN seamless Data Center Interconnect (DCI) with enhanced optimized intersubnet multicast (OISM). Without this feature, the DCI gateway devices flood multicast traffic across the interconnecting WAN. Flooding consumes significant WAN bandwidth if your network has many multicast flows or high multicast traffic rates. This feature seamlessly replaces the multicast flooding behavior. To optimize multicast forwarding across a DCI, OISM leaf devices in each network:
 - Advertise selective multicast Ethernet tag (SMET) routes (EVPN Type 6 routes) when a receiver subscribes to a multicast flow. The DCI gateways seamlessly propagate the SMET routes across the DCI on the OISM SBD.
 - Send multicast traffic based on the received SMET routes only to the remote receivers across the DCI who subscribed to that multicast flow.

To configure this feature:

- Configure the DCI gateway devices the same way you would configure the devices without multicast support.
- Configure enhanced OISM in the networks on both sides of the DCI.
 - With enhanced OISM, you can configure each OISM device with only the VLANs that the device hosts, except on multihoming peer OISM devices and the peer DCI gateways in each data center network. On those peer devices, you must configure the OISM revenue VLANs symmetrically.
 - Configure the same OISM supplemental bridge domain (SBD) in the matching tenant virtual routing and forwarding (VRF) instances on both sides of the DCI.

As OISM devices, the DCI gateways follow the enhanced OISM operational model to forward traffic to other OISM devices in their own network or across the DCI. They send the multicast traffic:

- Across the DCI to the other DCI gateways only on the OISM SBD.
- To other non-multihoming peer provider edge (PE) devices in their network only on the OISM SBD.

- To their multihoming peer PE devices only on the source VLAN.

You can also configure a DCI gateway as an OISM PIM EVPN gateway (PEG). The device acts as a DCI gateway and also as an OISM PEG border leaf device to exchange multicast traffic with devices outside of either network.

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

- **Generation of EVPN Type 3 routes on 802.1X dynamically mapped interfaces (EX4100-24MP, EX4300-MP, EX4400-24P, QFX5120-32C, QFX5120-48T, and QFX5120-48Y)**—Starting in Junos OS Release 24.2R1, we support generating Ethernet VPN (EVPN) Type 3 routes across interfaces dynamically mapped by the 802.1X protocol to a Virtual Extensible LAN (VXLAN) extended bridge domain (BD).

[See [dot1x](#).]

Junos OS API and Scripting

- **Support for configuring the `allow-transients` statement for individual commit scripts (EX4100-24MP, EX4400-24MP, and QFX5120-32C)**—Starting in Junos OS Release 24.2R1, you can configure the `allow-transients` statement for individual commit scripts. Configuring the `allow-transients` statement for individual scripts enables you to add transient configuration changes to the checkout configuration for specific commit scripts while still keeping transient changes disabled for the other commit scripts.

[See [allow-transients](#).]

Junos Telemetry Interface

- **OpenConfig MAC address and MAC address and IP path sensor support (ACX710, ACX5448, ACX5448-M, ACX5448-D, EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-48MP, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-24T, EX4100-F-12P, EX4100-F-48T, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, MX150, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vMX, QFX10002-60C, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, QFX10002, QFX10008, and QFX10016)**—Junos OS Release 24.2R1 supports the streaming of telemetry MAC address and MAC address and IP path data from the forwarding database to a collector using the OpenConfig resource path `/network-instances/network-instance/fdb/`. This feature is based on data models `openconfig-network-instance.yang` (version 1.2.0) and `openconfig-network-instance-l2.yang` (version 1.2.0).

[See [Junos YANG Data Model Explorer](#).]

- **Hardware resource threshold monitoring for capacity planning (EX4100-48MP, EX4400-24MP, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM0)**—Junos OS Release

24.2R1 supports hardware resource threshold monitoring. Using this feature, you can monitor hardware resource utilization. Advance knowledge about resource utilization nearing or crossing a maximum capacity provides time for you to act and prevent network disruption and traffic loss.

Use the `system packet-forwarding-options hw-resource-monitor resource-list` configuration statement at the `[edit]` hierarchy level to create a list of hardware resources that you want to monitor. Once configured, periodic resource monitoring occurs at the polling interval you set.

View the monitored data using operational mode commands or use Junos Telemetry interface (JTI) to send data from your device to a collector using the resource path `/junos/system/linecard/npu/memory/`.

[See [Configure Hardware Threshold Monitoring for Capacity Planning](#). For sensors, see [Junos YANG Data Model Explorer](#).]

MPLS

- **Support for IPv4 static route over IPv6 next-hop (MX204, MX240, MX304, MX480, MX960, MX10003, MX10016, MX2020, QFX5110, and QFX5200)**—Starting in Junos OS Release 24.2R1, you can configure an IPv4 static route over an IPv6 next hop to enable routing of IPv4 packets through the IPv6 next hop. The following IPv4 static route over IPv6 next-hop are supported:
 - IPv4 static route over IPv6 direct next-hop
 - IPv4 static route over IPv6 indirect-next-hop
 - IPv4 static route over IPv6 next-hop with preference

Use the following configuration to support IPv4 static route over IPv6 next-hop:

```
user@host# set routing-option static route ipv4-address next-hop ipv6-address
```

- **Provision binding SIDs for uncolored SR-TE (SR-MPLS) LSP (MX480 and QFX5200)**—Starting in Junos OS Release 24.2R1, we support provisioning of binding SID for uncolored SR-TE LSP where PCE requests PCC to allocate a binding SID from PCC's label space as follows:
 - PCE requests PCC to allocate a specific binding SID
 - PCE requests PCC to allocate binding SID of PCCs choice

We support the following PCE functionalities:

- PCE requests PCC to allocate binding SID of PCCs choice for delegated LSP.
- PCE requests PCC to allocate binding SID of PCCs choice for PCE-initiated LSP.
- PCE requests PCC to allocate a specific binding SID for delegated LSP.

- PCE requests PCC to allocate a specific binding SID for PCE-initiated LSP.
- Multiple candidate paths with binding SID in a policy.

We now support both 20-bit and 32-bit binding SID provisioned or requested from a PCE controller.

[See [PCEP Configuration](#).]

Network Management and Monitoring

- **AES-256 Encryption Algorithm Support for SNMPv3 (ACX5448, ACX5448-M, ACX5448-D, ACX710, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 24.2R1, you can configure Advanced Encryption Standard (AES) 256 algorithm for an SNMPv3 user. To configure AES-256 algorithm for an SNMPv3 user, include the `privacy-aes256` statement at the `[edit snmp v3 usm local-engine user username]` hierarchy level. AES-256 is a symmetric encryption algorithm that uses a 256-bit key to encrypt or decrypt messages and provides high-level security for protecting sensitive information.

[See [Configure SNMPv3 Encryption Type](#).]

- **Clear LLDP neighbors from an interface with the gRPC Network Operations Interface (gNOI) Layer2 service (ACX710, ACX5448, ACX5448-M, ACX5448-D, EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200 and QFX5210)**—Starting in Junos OS Release 24.2R1, you can execute supported Layer2 service remote procedure calls (RPCs) to perform the equivalent of the `clear lldp neighbors interface interface-name` command.

[See [gNOI Layer 2 Service](#).]

Routing Protocols

- **BGP link bandwidth community (cRPD, EX4100-48MP, EX4300-MP, EX4400-48MP, EX4650, EX9204, EX9208, MX240, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020, cSRX, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200, and QFX5210)**—Starting in Junos OS Release 24.2R1, BGP can communicate link speeds to remote peers, enabling better optimization of traffic distribution for load balancing. A BGP group can send the *link-bandwidth* non-transitive extended community over an EBGp session for originated or received and readvertised link-bandwidth extended communities.

To configure the non-transitive link bandwidth extended community, include the `bandwidth-non-transitive: value` in the export policy at the `[edit policy-options community name members community-ids]` hierarchy level.

To enable the device to automatically detect and attach the link-bandwidth community on a route at import, include the `auto-sense` auto-sense statement at the `[edit protocols bgp group link-bandwidth]` hierarchy level. This feature facilitates the integration of devices with different transmission speeds within the network, enabling efficient traffic distribution based on link speed.

[See and [group \(Protocols BGP\)](#).]

Software Installation and Upgrade

- **Base OS update (ACX710, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, Junos OS uses the FreeBSD main base OS. This upgrade provides improved security and better performance. In earlier releases, Junos OS used the FreeBSD Release 12 base OS.

[See [Junos® OS Software Installation and Upgrade Guide](#).]

- **In-band ZTP management in campus fabrics (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4100-24P, EX4100-24T, EX4100-F-48P, EX4100-F-24P, EX4100-F-48T, EX4100-F-24T, EX4100-F-12P, EX4100-F-12T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, EX9204, EX9208, EX9214, MX304, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 24.2R1, you can simplify the provisioning process for campus fabrics. Already provisioned upstream devices, such as core and distribution devices, that are capable of detecting downstream Day 0 devices can provide Layer 3 connectivity. With Layer 3 connectivity, the downstream Day 0 devices can proceed with Secure ZTP.

To configure in-band ZTP management, enable the `in-band-ztp` statement at the `[edit system services]` hierarchy on your core and distribution devices. Optionally, your cloud controller can provide the in-band-ztp configuration as part of the provisioning process for your core and distribution devices.

See [Zero Touch Provisioning](#)

- **Migration of Linux OS version**—Starting in Junos OS Release 24.2R1, the following devices support Wind River Linux LTS22:

Table 9: List of devices that support Wind River Linux LTS22

Platforms	Routing Engines Supported
ACX5448, ACX5448-D, and ACX5448-M	RE-ACX-5448

Table 9: List of devices that support Wind River Linux LTS22 (Continued)

Platforms	Routing Engines Supported
EX9204, EX9208, and EX9214	EX9200-RE2
MX240, MX480, and MX960	RE-S-X6
MX2010, MX2020	REMX2K-X8
MX2008	REMX2008-X8-64G
MX10008, MX10004	JNP10K-RE1
MX204	MX204
MX10003	JNP10003-RE1
MX304	JNP304-RE
SRX1600	SRX1600
SRX2300	SRX2300
SRX5800, SRX5600, and SRX5400	SRX5K-RE3
QFX10002-60C	RE-QFX10002-60C

Starting in Junos OS Release 24.2R1, in order to install VM Host image based on Linux WR LTS22, you have to upgrade the i40e NVM firmware version to 9.1 or later.

Additional Features

We have extended support for the following features to these platforms.

- **Supported transceivers, optical interfaces, and DAC cables** Select your product in the Hardware Compatibility Tool (<https://apps.juniper.net/hct/product/>) to view supported transceivers, optical

interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

- **Precision Time Protocol (PTP) enterprise profile** (QFX5110-48S)

[See [PTP Enterprise Profile](#).]

- **Wake-on LAN targeted-broadcast feature for EVPN-VXLAN networks** (EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4100-F-12P, EX4100-F-12T, EX4100-F-24P, EX4100-F-24T, EX4100-F-48P, EX4100-F-48T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)

[See [Targeted Broadcast](#) and [targeted-broadcast](#).]

What's Changed

IN THIS SECTION

- [EVPN | 145](#)
- [General Routing | 146](#)
- [Infrastructure | 147](#)
- [Routing Protocols | 147](#)
- [User Interface and Configuration | 147](#)
- [VPNs | 147](#)

Learn about what changed in this release for QFX Series Switches.



NOTE: For all QFX5110 models, the standard name of the image has been changed from “5e” to “5x.” As follows:

Old format: jinstall-host-qfx-5e-

New format: jinstall-host-qfx-5x-

The new format is in effect starting with Junos OS 24.2R1 and will be used for all subsequent mainline Junos OS releases. No maintenance or service releases for release trains prior to 24.2 will implement the change.

EVPN

- **OISM SBD bit in EVPN Type 3 route multicast flags extended community**—In EVPN Type 3 Inclusive Multicast Ethernet Tag (IMET) route advertisements for interfaces associated with the supplemental bridge domain (SBD) in an EVPN optimized intersubnet multicast (OISM) network, we now set the SBD bit in the multicast flags extended community. We set this bit for interoperability with other vendors, and to comply with the IETF draft standard for OISM, draft-ietf-bess-evpn-irb-mcast .

[See the description of the `show route table bgp.evpn.0` extensive command in [CLI Commands to Verify the OISM configuration](#).]

- **New commit check for MAC-VRF routing instances with the encapsulate-inner-vlan statement configured**— We introduced a new commit check that prevents you from configuring an IRB interface and the encapsulate-inner-vlan statement together in a MAC-VRF routing instance. Please correct or remove these configurations prior to upgrading to 23.2R2 or newer to avoid a configuration validation failure during the upgrade.

[See [encapsulate-inner-vlan](#).]

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the `[edit protocols evpn]` hierarchy level. In most use cases, you don't need to change the default limit. If you want to change the default limit, we recommend that you don't set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

[See [mac-ip-limit](#).]

- **Default behavior changes and new options for the easy EVPN LAG configuration (EZ-LAG) feature**—The easy EVPN LAG configuration feature now uses some new default or derived values, as follows:
 - Peer PE device `peer-id` value can only be 1 or 2.
 - You are required to configure the loopback subnet addresses for each peer PE device using the new `loopback peer1-subnet` and `loopback peer2-subnet` options at the `[edit services evpn device-attribute]` hierarchy level. The commit script uses these values for each peer PE device's loopback subnet

instead of deriving those values on each PE device. These replace the `loopback-subnet` option at the `[edit services evpn device-attribute]` hierarchy level, which has been deprecated.

- If you configure the `no-policy-and-routing-options-config` option, you must configure a policy statement called `EXPORT-L00` that the default underlay configuration requires, or configure the new `no-underlay-config` option and include your own underlay configuration.
- The commit script generates "notice" messages instead of "error" messages for configuration errors so you can better handle `[edit services evpn]` configuration issues.
- The commit script includes the element names you configure (such as IRB instance names and server names) in description statements in the generated configuration.

This feature also now includes a few new options so you have more flexibility to customize the generated configuration:

- `no-underlay-config` at the `[edit services evpn]` hierarchy level—To provide your own underlay peering configuration.
- `mtu overlay-mtu` and `mtu underlay-mtu` options at the `[edit services evpn global-parameters]` hierarchy level—To change the default assigned MTU size for underlay or overlay packets.

General Routing

- Starting in Junos OS Release 24.2R1, when you run the `run show lldp local-information interface <interface-name> | display xml` command, the output is displayed under the `lldp-local-info` root tag and in the `lldp-local-interface-info` container tag. When you run the `run show lldp local-information interface | display xml` command, the `lldp-tlv-filter` and `lldp-tlv-select` information are displayed under the `lldp-local-interface-info` container tag in the output.
- **Non-revertive switchover for sender based MoFRR**— In earlier Junos releases, source-based MoFRR ensured that the traffic reverted to the primary path from the backup path, when the primary path or session was restored. This reversion could result in traffic loss. Starting in Junos OS 22.4R3-S1, source-based MoFRR will not revert to the primary path, i.e. traffic will continue to flow through the backup path as long as the traffic flow rate on the backup path does not go below the configured threshold set under the `protocols mvpn hot-root-standby min-rate` command.
- Show active forwarding session for sender based MoFRR— The `show multicast route extensive` command will show the active forwarding session in the case of source-based MoFRR. The field `Session Status: Up & Forwarding` will indicate that the particular session is currently forwarding traffic.

Infrastructure

- **Option to disable path MTU discovery**—Path MTU discovery is enabled by default. To disable it for IPv4 traffic, you can configure the `no-path-mtu-discovery` statement at the `[edit system internet-options]` hierarchy level. To reenale it, use the `path-mtu-discovery` statement.

[See [Path MTU Discovery](#).]

Routing Protocols

- **Optimized mesh group routes (QFX5110, QFX5120, QFX5130, QFX5700 and ACX Series)**—The `show route snooping` command for `inet.1/inet6.1` table and `show route snooping table inet.1/inet6.1` will display only CE mesh group routes for platforms that support EVPN-MPLS or EVPN-VxLAN multicast. In earlier releases, other mesh groups like the VE mesh group were also displayed.

User Interface and Configuration

- **Configuration database maximum size increased (ACX Series, EX Series, MX Series, QFX Series, SRX Series, and vSRX)**—We've enhanced the `extend-size` statement at the `[edit system configuration-database]` hierarchy level to increase the maximum database size. On devices with a default configuration database size of ~400 MB, `extend-size` increases the maximum database size to ~2 GB. On devices with a default configuration database size of ~660 MB, `extend-size` increases the maximum database size to ~2.2 GB.

[See [configuration-database](#).]

VPNs

- **Increase in revert-delay timer range**— The `revert-delay` timer range is increased to 600 seconds from 20 seconds.

[See [min-rate](#).]

Known Limitations

There are no known limitations in hardware or software in this release for QFX Series switches.

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

Open Issues

IN THIS SECTION

- [General Routing](#) | 148
- [Routing Protocols](#) | 149

Learn about open issues in this release for QFX Series switches.

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

General Routing

- When TISSU upgrades from Junos OS Release 22.4 release and later, the box comes up as backup Routing Engine. [PR1703229](#)
- 4x25G channelized interfaces does not come up after swapping the optics hot. [PR1719758](#)
- An issue with SOFT OIR occurs, which is used for internal debugging purposes. [PR1757704](#)
- In a QFX51200-48YM-8C Virtual Chassis setup, after a mastership switch over fan tray of line card might not be displayed in the `show chassis hardware` and `show chassis environment` commands. There is no functional impact. [PR1758400](#)
- On QFX5210 devices, the `dcufe` process generates core file at `__kernel_vsyscall`, `tvp_watchdog`, and `dcbcm_driver_read32,soc_dcbcm_ipoll_check,cp u_sched_update_timers`. [PR1790234](#)

Routing Protocols

- On Junos platforms and Junos Evolved platforms, if a BGP peer goes down and stays down, the system might take an extremely long time to complete removing the BGP routes. The issue is observed when a BGP peer sends many routes, only a small number of routes are selected as the active routes in the routing information base (RIB, also known as the routing table), and if the BGP delete job gets only a small part of the CPU time because other work in the routing process is utilizing the CPU. [PR1695062](#)

Resolved Issues

IN THIS SECTION

- [General Routing | 149](#)
- [Network Management and Monitoring | 150](#)
- [Routing Protocols | 151](#)

Learn about the issues fixed in this release for QFX Series switches.

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

General Routing

- DHCP-relay does not work when no-snoop gets enabled on the QFX10002-60C devices. [PR1728639](#)
- On QFX10000 devices, proxy-arp restricted does not work as expected. [PR1769908](#)
- Family ethernet-switching policer per-sub-unit interface breaks after dcpfe/device restarts. [PR1771630](#)
- The dcpfe process crash due to stale memory. [PR1774366](#)
- Traffic drop occurs when VIPs become unreachable due to GARP sent on VLANs to which the VIP does not belong. [PR1778725](#)

- On QFX5110-48S devices, the virtual chassis port goes down when the VC is upgraded or the VC ports are deleted and added back. [PR1779624](#)
- Error gets generated on system when pvidb variable gets accessed. [PR1781317](#)
- The fxpc process crashes and the device reboots after deleting Aggregated Ethernet (AE) Interface along with its associated physical interface and then applying new interface configuration on the associated physical interface in an EVPN-VXLAN scenario. [PR1783397](#)
- Traffic loss after PIC restart if the packet has a VLAN tag of 4095. [PR1788573](#)
- Traffic loss observed with line rate jumbo frames. [PR1789302](#)
- Packet corruption on QFX10000 devices occurs in the egress pipeline. [PR1792732](#)
- Nexthop resolution fails in a high scale ARP on QFX10000 devices. [PR1793335](#)
- FPC crash (dcpfe process core-dump) and traffic drop occurred due to "Err Detect Register" and "Data Cache Parity Error". [PR1795339](#)
- In the EVPN-VXLAN scenario, traffic blackholes on spine reboot. [PR1796210](#)
- The dcpfe process crashes while enabling loop-detect on the SP style interface. [PR1796883](#)
- On QFX-10002-36q devices, the SFP+-10G-CU3M SFP does not come up post upgrade from Junos OS Release 18.2X75-D12.6 to Junos OS Release 20.4R3-S7.2. [PR1797453](#)
- The dcpfe process generates core file on QFX10000 devices while running profile baseline in Junos OS Release 22.2R3-S3.18 image. [PR1797511](#)
- DCPFE process crashes in an EVPN-VXLAN scenario. [PR1797516](#)
- On QFX5000 devices, the ECMP programming issue occurs in EVPN-VXLAN. [PR1802958](#)
- Traffic loss occurs when the layer 3 interface gets deleted. [PR1808550](#)

Network Management and Monitoring

- Warning messages gets generated while Custom Yang Package Deletion. [PR1804856](#)

Routing Protocols

- The rpd process crashes when the routing-instances name length is greater than 60 characters.
[PR1795964](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrading Software on QFX Series Switches | 151](#)
- [Installing the Software on QFX10002-60C Switches | 153](#)
- [Installing the Software on QFX10002 Switches | 154](#)
- [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 | 155](#)
- [Installing the Software on QFX10008 and QFX10016 Switches | 156](#)
- [Performing a Unified ISSU | 160](#)
- [Preparing the Switch for Software Installation | 161](#)
- [Upgrading the Software Using Unified ISSU | 161](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 164](#)

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

Upgrading Software on QFX Series Switches

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



NOTE: For all QFX5110 models, the standard name of the image has been changed from “5e” to “5x.” As follows:

Old format: jinstall-host-qfx-5e-

New format: jinstall-host-qfx-5x-

The new format is in effect starting with Junos OS 24.2R1 and will be used for all subsequent mainline Junos OS releases. No maintenance or service releases for release trains prior to 24.2 will implement the change.

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 **24.2** 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 24.2 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-x86-64-24.2-R1.n-secure-signed.tgz reboot
```

Replace *source* with one of the following values:

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

Adding the `reboot` command reboots the switch after the upgrade is installed. When the reboot is complete, the switch displays the login prompt. The loading process can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 24.2 `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-x86-64-20.4R1.9.tgz
```

If the Install Package resides remotely from the switch, execute the **request vmhost software add <pathname><source>** command.

For example:

```
user@switch> request vmhost software add ftp://ftpserver/directory/junos-vmhost-install-qfx-x86-64-20.4R1.9.tgz
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10002 Switches



NOTE: If you are upgrading from a version of software that does not have the FreeBSD 10 kernel (15.1X53-D30, for example), you will need to upgrade from Junos OS Release 15.1X53-D30 to Junos OS Release 15.1X53-D32. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D60 or Junos OS Release 18.3R1.



NOTE: On the switch, use the **force-host** option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the **force-host** option.

If the installation package resides locally on the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-f-x86-64-20.4R1.n-secure-signed.tgz reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-20.4R1.n-secure-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches



NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add <pathname><source>** command.

To install the software on re0:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-
m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-
domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-
m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.



NOTE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.



WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI `delete chassis redundancy` command when prompted. If GRES is enabled, it will be removed with the redundancy command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the `[edit routing-options]` hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

```
user@switch# delete chassis redundancy
```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the `request system software add` command:

```
user@switch> request system software add validate /var/tmp/jinstall-host-qfx-10-f-
x86-64-20.4R1.n-secure-signed.tgz
```

For more information about the `request system software add` command, see the [CLI Explorer](#).

9. Reboot the switch to start the new software using the `request system reboot` command:

```
user@switch> request system reboot
```



NOTE: You must reboot the switch to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the `request system software delete <package-name>` command. This is your last chance to stop the installation.

All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the `show version` command to verify the version of the software installed.

```
user@switch> show version
```

Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software.

11. Log in to the master Routing Engine console port.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the `request chassis routing-engine master` command, see the [CLI Explorer](#).

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

```
user@switch> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state          Backup
    Election priority      Master (default)

Routing Engine status:
  Slot 1:
    Current state          Master
    Election priority      Backup (default)
```

14. Install the new software package using the `request system software add` command:

```
user@switch> request system software add validate /var/tmp/jinstall-host-qfx-10-f-
x86-64-20.4R1.n-secure-signed.tgz
```

For more information about the `request system software add` command, see the [CLI Explorer](#).

15. Reboot the Routing Engine using the `request system reboot` command:

```
user@switch> request system reboot
```



NOTE: You must reboot to load the new installation of Junos OS on the switch. To abort the installation, do not reboot your system. Instead, finish the installation and then issue the `request system software delete jinstall <package-name>` command. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not send traffic.

16. Log in and issue the `show version` command to verify the version of the software installed.
17. Transfer routing control back to the master Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the `request chassis routing-engine master` command, see the [CLI Explorer](#).

18. Verify that the master Routing Engine (slot 0) is indeed the master Routing Engine:

```
user@switch> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state           Master
    Election priority       Master (default)

Routing 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 161](#)
- ["Upgrading the Software Using Unified ISSU" on page 161](#)

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-x86-64-20.4R1.n-secure-signed.tgz`.



NOTE: During the upgrade, you cannot access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-host-qfx-5-f-x86-64-18.3R1.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-19.2R1.n-secure-signed.tgz completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
```

```

ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```



NOTE: A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).



NOTE: If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrade and Downgrade Support Policy for Junos OS Releases

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 sixty months after the first general availability date and customer support for an additional six more months.



NOTE: The sixty months of support for EEOL releases is introduced in Junos OS 23.2 release and is available for all later releases. For releases prior to 23.2, the support for EEOL releases continues to be thirty six 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 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
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	60 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 | 165](#)
- [What's Changed | 165](#)
- [Known Limitations | 165](#)
- [Open Issues | 165](#)
- [Resolved Issues | 166](#)

What's New

There are no new features or enhancements to existing features in this release for Juniper Secure Connect.

What's Changed

There are no changes in behavior and syntax in this release for Juniper Secure Connect.

Known Limitations

There are no known limitations in hardware or software in this release 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 or software in this release 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 this release 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.

Junos OS Release Notes for SRX Series Firewalls

IN THIS SECTION

- [What's New | 166](#)
- [What's Changed | 187](#)
- [Known Limitations | 192](#)
- [Open Issues | 193](#)
- [Resolved Issues | 195](#)
- [Migration, Upgrade, and Downgrade Instructions | 197](#)
- [Documentation Updates | 199](#)

What's New

IN THIS SECTION

- [Hardware | 168](#)
- [Application Identification \(AppID\) | 178](#)
- [Chassis | 180](#)

- [Flow-Based and Packet-Based Processing | 180](#)
- [High Availability | 181](#)
- [Interfaces | 181](#)
- [Juniper Advanced Threat Prevention Cloud \(ATP Cloud\) | 182](#)
- [Juniper Extension Toolkit \(JET\) | 182](#)
- [J-Web | 182](#)
- [Network Management and Monitoring | 183](#)
- [Public Key Infrastructure \(PKI\) | 183](#)
- [Serviceability | 184](#)
- [Software Installation and Upgrade | 184](#)
- [VPNs | 185](#)
- [Additional Features | 186](#)

Learn about new features introduced in this release for SRX Series Firewall devices.

To view features supported on the SRX Series Firewall, view the Feature Explorer using the following links. To see which features were added in Junos OS Release 24.2R1, click the Group by Release link. You can collapse and expand the list as needed.

- [SRX300](#)
- [SRX320](#)
- [SRX340](#)
- [SRX345](#)
- [SRX380](#)
- [SRX1500](#)
- [SRX1600](#)
- [SRX2300](#)
- [SRX4100](#)
- [SRX4200](#)
- [SRX4300](#)

- [SRX4600](#)
- [SRX5400](#)
- [SRX5600](#)
- [SRX5800](#)

Hardware

- **New SRX4300 Firewall**—Starting in Junos OS Release 24.2R1, we introduce the midrange SRX4300 Firewall. The SRX4300 Firewall provides next-generation firewall capabilities and advanced threat detection and mitigation. This firewall is ideal for small and medium sized enterprise edge, campus edge, data center edge firewall, and secure VPN router deployments for distributed enterprise use cases.

Table 11: Features Supported on SRX4300 Firewall

Feature	Description
Chassis	<ul style="list-style-type: none"> • Chassis and field-replaceable unit (FRU) management support, including: <ul style="list-style-type: none"> • Temperature threshold monitoring using sensors • Power supply unit (PSU) control • PIC detection • Fabric management • Fan speed adjustment as per EM policy <p>[See Configuring Ambient Temperature and Chassis-Level User Guide.]</p>
Chassis Cluster	<ul style="list-style-type: none"> • Support for in-service software upgrade (ISSU) and dual control links with Media Access Control Security (MACsec) <p>[See Upgrading a Chassis Cluster Using In-Service Software Upgrade and Media Access Control Security (MACsec) on Chassis Cluster.]</p>

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none"> • Support for CoS <p>[See Understanding Class of Service.]</p>
Hardware	<ul style="list-style-type: none"> • The SRX4300 is a 1-U chassis with the following ports: <ul style="list-style-type: none"> • Eight 10 multi-rate Gigabit Ethernet interface (mge) BASE-T ports • Eight 10-Gigabit Ethernet (GbE) SFP+ ports • Four 25GbE SFP28 ports • Six 100GbE QSFP28 ports • Two 1GbE SFP HA ports <p>All ports are MACsec capable and support both AC and DC variants.</p> <p>To install the SRX4300 hardware and perform initial software configuration, routine maintenance, and troubleshooting, see SRX4300 Firewall Hardware Guide.</p> <p>[See Feature Explorer for the complete list of features for any platform.]</p>

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
High availability (HA) and resiliency	<ul style="list-style-type: none"> • Support for BFD <ul style="list-style-type: none"> • Support up to 3 x 300-millisecond (msec) failure detection time • Support up to 100 BFD sessions <p>[See Understanding BFD for Static Routes for Faster Network Failure Detection and Understanding How BFD Detects Network Failures.]</p> • Multinode High Availability supports Auto Discovery VPN (ADVPN) in node-local tunnel deployment. <p>Node-local tunnels enhance Multinode High Availability by providing separate tunnels from a VPN peer device to both nodes in the setup. With ADVPN, VPN tunnels can be established dynamically between spokes. Combining ADVPN with Multinode High Availability in node-local tunnel deployment ensures robust network connectivity, efficient resource utilization, and seamless failover capability.</p> <p>[See IPsec VPN Support in Multinode High Availability.]</p> • Support for Multinode High Availability in routing, hybrid, and default gateway modes <p>[See Multinode High Availability.]</p> • Provides platform software resiliency support for the following hardware components: <ul style="list-style-type: none"> • CPU • Peripheral Component Interconnect (PCI)

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • Memory • Solid state device (SSD) • Inter-integrated circuit (I2C) • Temperature sensor • Voltage sensor • Fan • Power supply units (PSUs) in 1+1 redundancy mode <p>When a hardware component fails, the Junos OS software:</p> <ul style="list-style-type: none"> • Logs the message with failure details, including time stamp, module name, and component name. • Raises or clears alarms, if applicable. • Makes the LED glow to indicate FRU fault. • Performs local action, such as self-healing and taking the component out of service. <p>[See Chassis-Level User Guide.]</p>

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
Interfaces	<ul style="list-style-type: none"> • Interfaces support includes four PICs with the following default speeds: <ul style="list-style-type: none"> • PIC 0 with 10 Gbps (Copper) • PIC 1 with 10 Gbps (SFP+) • PIC 2 with 25 Gbps (SFP28) • PIC 3 with 100 Gbps (QSFP28) <p>Junos OS creates PIC 0 by default. You can create PIC 1, PIC 2, and PIC 3 interfaces by inserting SFP+, SFP28, and QSFP28 transceivers, respectively.</p> <p>[See Port Speed on SRX Series Firewalls.]</p> <ul style="list-style-type: none"> • Mixed speed support on SFP28 ports. <p>You can configure two options in PIC mode; 1GbE/10GbE combined and 25GbE.</p> <p>[See Port Speed on SRX Series Firewalls.]</p>
Junos telemetry interface (JTI)	<ul style="list-style-type: none"> • Stream data from a device to a collector using basic JTI sensors and new flow monitoring sensors. Junos OS supports the following flow sensors: <ul style="list-style-type: none"> • PIC CPU utilization <code>/junos/security/spu/cpu</code> • Flow session and flow packets <code>/junos/security/spu/flow</code> • Flow session and flow packets for logical systems <code>/junos/security/spu/flow/lsys</code> <p>[For state sensors, see Junos YANG Data Model Explorer.]</p>

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
Layer 7 security features	<ul style="list-style-type: none"> • Support for advanced policy-based routing (APBR) [See Advanced Policy-Based Routing.] • Support for application identification (AppID) [See Application Identification.] • Support for application quality of experience (AppQoE) [See Application Quality of Experience.] • Support for application quality of service (AppQoS) [See Application QoS.] • Support for Content Security [See Content Security Overview.] • Support for intrusion detection and prevention (IDP) [See Intrusion Detection and Prevention Overview.] • Support for Juniper ATP Cloud [See File Scanning Limits.] • Support for Juniper Networks Deep Packet Inspection-Decoder (JDPI) [See Juniper Networks Deep Packet Inspection-Decoder (JDPI-Decoder).] • Support for SSL proxy [See SSL Proxy.]

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
MACsec	<ul style="list-style-type: none"> • Support for MACsec in static CAK mode on physical interfaces with the following encryptions: <ul style="list-style-type: none"> • GCM-AES-128 • GCM-AES-256 • GCM-AES-XPB-128 • GCM-AES-XPB-256 <p>Channelized ports and switch-to-switch connections support this feature.</p> <p>[See Configuring MACsec.]</p>
Network management and monitoring	<ul style="list-style-type: none"> • Support for filter-based packet capture for real-time data packets traveling over the network. Support for datapath debugging is not yet available. <p>[See Example: Configure a Firewall Filter for Packet Capture.]</p>
Remote access	<ul style="list-style-type: none"> • Support for remote access VPN using Juniper Secure Connect <p>[See Juniper Secure Connect Administrator Guide.]</p>

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
Services applications	<ul style="list-style-type: none"> • Support for Application Layer Gateway (ALG) [See ALG Overview.] • Support for ADVPN configuration with IPv6 address on firewalls that run the ikev2 process for IPsec VPN service [See Auto Discovery VPNs.] • Support for ChaCha20-Poly1305 authenticated encryption algorithm for IPsec VPN services [See proposal (Security IKE) and proposal (Security IPsec).] • Support for multicast traffic in AutoVPN and ADVPN with ikev2 process using PIM sparse mode over st0 P2MP interface on firewalls that run the ikev2 process for IPsec VPN service. Supports IPv4 multicast in PIM sparse mode. [See AutoVPN and Auto Discovery VPNs.] • Support for DNS [See Understanding and Configuring DNS, DNS ALG, DNS Proxy Overview, DNS Names in Address Books, and DNSSEC Overview.] • Support for user authentication [See User Authentication Overview.] • Support for security policies [See Configuring Security Policies.] • Support for security zones [See Security Zones.] • Support for Network Address Translation (NAT)

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
	<p data-bbox="898 359 1268 386">[See NAT Configuration Overview.]</p> <ul style="list-style-type: none"> <li data-bbox="862 422 1398 485">• Support for screens options for attack detection and prevention <p data-bbox="898 520 1386 583">[See Screens Options for Attack Detection and Prevention.]</p> <ul style="list-style-type: none"> <li data-bbox="862 619 1203 646">• Support for traffic processing <p data-bbox="898 682 1382 745">[See Traffic Processing on SRX Series Firewalls Overview.]</p> <ul style="list-style-type: none"> <li data-bbox="862 781 1263 808">• Support for integrated user firewall <p data-bbox="898 844 1321 871">[See Configure Integrated User Firewall.]</p> <ul style="list-style-type: none"> <li data-bbox="862 907 1414 1003">• Support for IPsec VPN with ike process. Support for the policy-based VPN and Group VPN is not yet available. <p data-bbox="898 1018 1328 1045">[See IPsec VPN Configuration Overview.]</p> <ul style="list-style-type: none"> <li data-bbox="862 1081 1274 1108">• Support for PowerMode IPsec (PMI) <p data-bbox="898 1144 1149 1171">[See PowerMode IPsec.]</p> <ul style="list-style-type: none"> <li data-bbox="862 1207 1089 1234">• Support for DHCP <p data-bbox="898 1270 1133 1297">[See DHCP Overview.]</p> <ul style="list-style-type: none"> <li data-bbox="862 1333 1177 1360">• Support for GTP and SCTP <p data-bbox="898 1396 1414 1423">[See Monitoring GTP Traffic and SCTP Overview.]</p> <ul style="list-style-type: none"> <li data-bbox="862 1459 1198 1486">• Support for on-box reporting <p data-bbox="898 1522 1170 1549">[See report (Security Log).]</p> <ul style="list-style-type: none"> <li data-bbox="862 1585 1321 1612">• Support for inline active flow monitoring <p data-bbox="898 1648 1398 1675">[See Understand Inline Active Flow Monitoring.]</p> <ul style="list-style-type: none"> <li data-bbox="862 1711 1112 1738">• Support for TWAMP

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
	<p>[See Understand Two-Way Active Measurement Protocol.]</p> <ul style="list-style-type: none"> • Support for RPM <p>[See Real-Time Performance Monitoring for SRX Devices.]</p> <ul style="list-style-type: none"> • Support for logical systems <p>[See Logical Systems Overview.]</p>
Software Installation and Upgrade	<ul style="list-style-type: none"> • Support for BIOS, secure boot, and bootloader <p>[See Secure Boot and Bootloader]</p> <ul style="list-style-type: none"> • Support for jfirmware <p>[See Installing and Upgrading Firmware, request system firmware upgrade, and show system firmware.]</p> <ul style="list-style-type: none"> • Support for secure zero-touch provisioning (ZTP) <p>[See Secure Zero Touch Provisioning.]</p>

Table 11: Features Supported on SRX4300 Firewall *(Continued)*

Feature	Description
User access and authentication administration	<p>Support for Trusted Platform Module (TPM)-based certificates for advanced anti-malware (AAMW) protection To use the TPM-based certificates:</p> <ul style="list-style-type: none"> • The device loads the TPM-based certificate using PKI during the device's start and restart operations. To view the TPM-based certificate ID, referred to as <code>idev-id</code>, use the <code>show security pki node-local local-certificate certificate-id idev-id</code> command. • The SSL Initiation uses the certificate for Transport Layer Security (TLS) connection to authenticate the device. You can configure the <code>tpm</code> option using the <code>set services ssl initiation profile profile-name crypto-hardware-offload</code> command. <p>See show security pki node-local local-certificate and profile (SSL Initiation).]</p>

Application Identification (AppID)

- **Application signature package installation enhancements (SRX Series Firewalls and vSRX)**—Starting in Junos OS Release 24.2R1, we've enhanced application signature package installation with the following changes:
 - During application signature package installation, the system performs data plane validation. This validation checks for errors in the package. If successful, the installation proceeds. If errors are found, the installation stops and reverts to the previous active version.
 - When using a chassis cluster setup, the system first installs the application signature package on the primary node and checks for any issues or problems. If the validation is successful, it then proceeds to install the same package on the secondary node.
 - The auto rollback feature now enables the system to revert to a previously working version of the application signature package. Additionally, it retains the previously designated rollback version in the event of any issues during application signature package installation.

New enhancements ensure a smooth transition by reverting to a known working version if needed.

See [[Application Signatures for Application Identification](#)].

- **CASB support (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos Release OS 24.2R1, SRX Series Firewalls support Cloud Access Security Broker (CASB).

CASB discovers SaaS applications in use and provides visibility and granular controls to protect and manage access to cloud applications. On SRX Series Firewalls, CASB provides inline activity control for the following set of cloud applications:

- Box
- Dropbox
- Salesforce
- Google Docs
- OneDrive
- SharePoint
- Slack
- Gmail

See [[Cloud Access Security Broker \(CASB\) Policy](#)].

- **SSL proxy enhancements (SRX Series Firewalls and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we introduce following enhancements for SSL proxy on SRX Series Firewalls:
 - Support of SNI extension at SSL initiation (SSL-I).
 - Support of certificate chain at SSL-I for client certificate verification.
 - Support for P-384, P-512 EC group for SSL proxy profile in addition to P-256.
 - Support for new ECDSA ciphers for SSL initiation and SSL termination profiles in non-proxy mode:
 - ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
 - ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
 - ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
 - ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
 - ECDHE_ECDSA_WITH_AES_256_CBC_SHA
 - ECDHE_ECDSA_WITH_AES_128_CBC_SHA

- ECDHE_ECDSA_WITH_CHACHA20_POLY1305
- New syslog messages for SSL configurations.
 - **SSL_CONFIG_MEMORY_ALLOCATION_FAILURE**— For memory allocation
 - **SSL_CONFIG_PROFILE_PROCESS_ERR** —For SSL profile processing
 - **SSL_CONFIG_CERT_PROCESS_ERR**— For SSL certificate processing.
 - **SSL_GLOBAL_CONFIG_PROCESS_ERR**—For SSL global configuration.
 - **SSL_CONFIG_PKI_IPC_ERR**— For IPC communication for SSL-PKI

See [[Cipher Suites for SSL Proxy.](#)]

Chassis

- **Resiliency support (SRX2300)**—Starting in Junos OS Release 24.2R1, the resiliency feature is enabled for the following hardware components:
 - CPU
 - Peripheral Component Interconnect (PCI)
 - Memory
 - Inter-Integrated Circuit (I2C)
 - Temperature sensor
 - Two power supply units (PSUs) in 1+1 redundancy mode
 - Fan

Resiliency feature monitors the platform components periodically, performs alarm management, and takes corrective actions if an anomaly is persistently encountered.

[See [Chassis-Level User Guide.](#)]

Flow-Based and Packet-Based Processing

- **Decouple inet and mpls (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, an SRX Series Firewall working in packet mode does not forward traffic anymore after the Junos OS upgrade. You must configure set security forwarding-options family inet mode packet-based immediately after the Junos upgrade to restore the operation of the device in packet mode.

The inet family, which was coupled with the mpls family prior to Junos OS Release 24.2R1, is now decoupled from the mpls family. You can enable packet mode for the inet family separately.

[See [Packet-Based Forwarding](#).]

High Availability

- **ADVPN support on node-local tunnels in Multinode High Availability (SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we support Auto Discovery VPN (ADVPN) on node-local tunnels configured with Multinode High Availability.

Node-local tunnels enhance Multinode High Availability by providing separate tunnels from a VPN peer device to both nodes in the setup. With ADVPN, VPN tunnels can be established dynamically between spokes. Combining ADVPN with Multinode High Availability in a node-local tunnel deployment ensures robust network connectivity, efficient resource utilization, and seamless failover capability.

See [\[VPN Support in Multinode High Availability\]](#).

- **Features support for asymmetric traffic flows in Multinode High Availability (SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 24.2R1, a Multinode High Availability setup supports the following features for asymmetric traffic flows:
 - Network Address Translation (NAT)
 - Carrier-Grade NAT (CGNAT)
 - Application Layer Gateway (ALG)
 - GPRS Tunneling Protocol (GTP)
 - User firewall and firewall authentication
 - Layer 7 services (intrusion detection and prevention (IDP), application identification (AppID), Content Security, application quality of service (AppQoS), advanced policy-based routing (APBR), unified policy, user firewall authentication, onbox-logging, and SSL proxy)

See [\[Asymmetric Traffic Flow Support for Multinode High Availability\]](#) and [\["Known Limitations" on page 192\]](#).

Interfaces

- **Mixed speed support on SFP28 ports (SRX1600 and SRX2300)**—Starting in Junos OS Release 24.2R1, you have two configuration options in PIC mode: 1GbE/10GbE combined and 25GbE.

[See [Port Speed on SRX Series Firewalls](#).]

Juniper Advanced Threat Prevention Cloud (ATP Cloud)

AI-Predictive Threat Prevention leverages machine learning-based zero-day threat detection (SRX Series Firewall and vSRX Series Firewall)—Starting in Junos OS Release 24.2R1, you can configure machine learning-based threat detection for zero-day threats at line rate. File scanning during threat detection happens without Internet access and only a small section of file data is sufficient for the detection to return a verdict. Machine learning-based threat detection becomes available on your firewall when the latest antivirus signature pack is automatically downloaded from the Juniper Networks content delivery network (CDN) server to your firewall.

[See [Example: Configure Flow-Based Antivirus Policy](#), [anti-virus](#), and [show services anti-virus statistics](#).]

- **System log messages for GeoIP (SRX Series Firewalls and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we've enhanced the IP-based geolocation (GeoIP) feature to provide improved consistency checks and logging from SRX Series Firewalls that are enrolled with Juniper ATP Cloud.

The session deny message includes the following fields:

- **source-country**—Displays the country code of the source address with reference to the policy dynamic address match.
- **destination-country**—Displays the country code of the destination address with reference to the policy dynamic address match.

The system log message displays the valid country code only if the matched policy includes a dynamic address configured with GeoIP. If the matched policy does not have GeoIP configured, then the source-country and destination-country fields display N/A.

[See [System Log Explorer](#) and [Configure Juniper Advanced Threat Prevention Cloud With Geolocation IP](#).]

Juniper Extension Toolkit (JET)

- **JET MAP-E Service API (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, you can use the JET MAP-E Service API to remotely configure Mapping of Address and Port with Encapsulation (MAP-E) rules in your network.

[See [Overview of JET APIs](#) and [Mapping of Address and Port with Encapsulation \(MAP-E\)](#).]

J-Web

- **Support for SRX4300 Firewall (SRX4300)**— Starting in Junos OS Release 24.2R1, J-Web supports SRX4300 Firewall.

[See [The J-Web Setup Wizard](#), [Dashboard Overview](#), [Monitor Interfaces](#), and [About Reports Page](#).]

Network Management and Monitoring

- **Logging Infrastructure Support for RADIUS Accounting (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we've introduced a logging infrastructure for RADIUS accounting. The RADIUS accounting enables you to send the logs directly from dataplane to the RADIUS accounting server. When you enable RADIUS logging, logs are sent for NAT PBA ALLOC, INTERIM, and RELEASE events to the configured RADIUS server. This feature enhances the existing stream-based logging and includes:
 - Incorporation of vendor-specific attributes (VSAs) in RADIUS accounting messages
 - Support multiple RADIUS accounting servers under different streams
 - Manage retries and retransmissions of RADIUS accounting messages in case of failure
 - Flexible and capable of supporting a backup RADIUS accounting server.

To support the feature, we've introduced the following configuration statements:

- radius
- retry-count
- radius-accounting
- subscriber-extension

Use the following commands to view and to clear the RADIUS server counters for RADIUS streams:

- show security log radius stream
- clear security log radius stream.

[See [radius \(Security Log\)](#), [retry-count \(Security Log\)](#), [radius-accounting](#), and [subscriber-extension](#).]

Public Key Infrastructure (PKI)

- **TPM-based certificate support for advanced anti-malware protection (SRX1600, SRX2300, and SRX4300)**—Starting in Junos OS Release 24.2R1, your SRX Series Firewalls use Trusted Platform Module based (TPM-based) certificates for advanced anti-malware (AAMW) protection. To use the TPM-based certificates:
 - The firewall loads the certificate using PKI during the device's start and restart operation. To view the certificate ID, referred as *idev-id*, use the show security pki node-local local-certificate certificate-id idev-id command.

- The SSL Initiation uses the certificate for Transport Layer Security (TLS) connection to authenticate the device. You can configure the `tpm` option using the `set services ssl initiation profile profile-name crypto-hardware-offload` command.

[See [show security pki node-local local-certificate](#), and [profile \(SSL Initiation\)](#).]

Serviceability

- **Datapath trace debug (SRX Series Firewalls)**—Starting in Junos OS Release 24.2R1, we've enhanced the datapath trace debug. With this enhancement, the flow trace provides:
 - Clear and concise information
 - Packet tracking information
 - Information about missing trace logs
 - Ensure to capture a finite number of initial trace messages

[See [Understanding Data Path Debugging for SRX Series Devices](#).]

- **IDP signature package improvements (SRX Series Firewalls and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we've made enhancements to improve the reliability of package updates and validation of update installation. The system automatically rolls back the signature package if the security package installation fails.

In the case of a multi-SPC/PIC device, when a failure occurs in the data plane after a signature pack is installed, the security package is rolled back once the PICs come back online. This action limits potential damage.

In a high-availability environment, to prevent failovers from occurring in a loop, the integrity validation is conducted during the primary node installation stage. The system installs the security package on the secondary node only after confirming that the package has no integrity issues.

[See [IDP Signature Database Overview](#).]

Software Installation and Upgrade

- **Base OS update (ACX710, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, Junos OS uses the FreeBSD main base OS. This upgrade provides improved security and better performance. In earlier releases, Junos OS used the FreeBSD Release 12 base OS.

[See [Junos® OS Software Installation and Upgrade Guide](#).]

- **Migration of Linux OS version**—Starting in Junos OS Release 24.2R1, the following devices support Wind River Linux LTS22:

Table 12: List of devices that support Wind River Linux LTS22

Platforms	Routing Engines Supported
ACX5448, ACX5448-D, and ACX5448-M	RE-ACX-5448
EX9204, EX9208, and EX9214	EX9200-RE2
MX240, MX480, and MX960	RE-S-X6
MX2010, MX2020	REMX2K-X8
MX2008	REMX2008-X8-64G
MX10008, MX10004	JNP10K-RE1
MX204	MX204
MX10003	JNP10003-RE1
MX304	JNP304-RE
SRX1600	SRX1600
SRX2300	SRX2300
SRX5800, SRX5600, and SRX5400	SRX5K-RE3
QFX10002-60C	RE-QFX10002-60C

Starting in Junos OS Release 24.2R1, in order to install VM Host image based on Linux WR LTS22, you have to upgrade the i40e NVM firmware version to 9.1 or later.

VPNs

- **Remote access VPN support (SRX2300)**—Starting in Junos OS Release 24.2R1, you can use Juniper Secure Connect for remote access VPN. [See [Juniper Secure Connect Administrator Guide](#).]

- **Support for ChaCha20-Poly1305 algorithm (SRX1600, SRX2300, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—Starting in Junos OS Release 24.2R1, we support ChaCha20-Poly1305 authenticated encryption algorithm for IPsec VPN services. You can configure the algorithm using the option `chacha20-poly1305` for:

- control plane with the IKEv2 protocol.
- data plane with the IPsec ESP protocol. You configure the algorithm in PowerMode IPsec (PMI) mode for the SRX Series Firewalls, and in both the PMI and non-PMI modes for vSRX 3.0. You cannot use the algorithm for IPsec when the VPN monitoring feature is enabled.

[See [proposal \(Security IKE\)](#), [proposal \(Security IPsec\)](#), [show security ike security-associations](#), and [show security ipsec security-associations](#).]

- **Support for IPv6 address in ADVPN with iked process (SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—Starting in Junos OS Release 24.2R1, we support Auto Discovery VPN (ADVPN) configuration with IPv6 address on firewalls that run the iked process for IPsec VPN service.

[See [Auto Discovery VPNs](#).]

- **Support for multicast traffic in AutoVPN and ADVPN with iked process (SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, and vSRX 3.0)**—Starting in Junos OS Release 24.2R1, we support IP multicast with AutoVPN and Auto Discovery VPN (ADVPN). The IP multicast uses Protocol Independent Multicast (PIM) using point-to-multipoint (P2MP) mode over st0 interface on firewalls that run the iked process for IPsec VPN service. Your firewall supports IPv4 multicast in PIM sparse mode.

You can enable PIM on the st0 secure tunnel interface using the *interface-name* option at the `[edit protocols pim interface interface-name]` hierarchy level.

[See [AutoVPN](#), [Auto Discovery VPNs](#), and [interface \(Protocols PIM\)](#).]

Additional Features

We have extended support for the following features to these platforms.

- **Support for preventing event, op, SNMP, or JET script execution based on current system memory usage** (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800)

[See [Configure Script Start Options](#).]

- **Supported transceivers, optical interfaces, and DAC cables** Select your product in the Hardware Compatibility Tool (<https://apps.juniper.net/hct/product/>) to view supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

- **On-box monitoring support on the control plane** (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800). You can configure the memory monitoring system to monitor the system memory and raise a major or minor alarm using the `set system monitor memory system alarm` command statement on a device. The system raises an alarm when the device runs low on memory.

[See [Memory \(System\)](#).]

- **Reduce operational time by using the JET Interfaces Service API to perform port bounces** (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800)

[See [Overview of JET APIs](#).]

What's Changed

IN THIS SECTION

- [Application Security | 187](#)
- [Authentication and Access Control | 188](#)
- [EVPN | 188](#)
- [Interfaces | 189](#)
- [Junos OS API and Scripting | 189](#)
- [User Interface and Configuration | 190](#)
- [VPNs | 190](#)

Learn about what changed in this release for SRX Series Firewalls.

Application Security

- **Application Signatures Package (SRX Series Firewalls and vSRX)**—The `show services application-identification status` command output displayed incorrect date for application package version release date. The command output displays the release date of the initial installed application signature package. Subsequent installations of newer versions do not update the release date of the signature

package. The release date is only updated correctly when installing a signature package that has changes in PB version/Engine version compared to the currently installed ones.

Starting in Junos OS Release 24.2 onwards, the command output shows the correct date.

See [show services application-identification status](#).

- **Deprecation of 3DES-CBC ciphers (SRX Series Firewalls and vSRX)**—Support for the following ciphers is deprecated:
 - RSA-3DES-EDE-CBC-SHA
 - ECDHE-ECDSA-3DES-EDE-CBC-SHA

The options to configure these ciphers are not available at the `[edit system services ssh]` hierarchy.

- Starting from Junos 21.4R1 platforms with the following Routing Engines which have Intel CPUs with microcode version 0x35 observe the error warning, "000: **Firmware Bug:** TSC_DEADLINE disabled due to Errata; please update microcode to version: 0x3a (or later)" on the console. RE-S-X6-64G RE-S-X6-128G REMX2K-X8-64G RE-PTX-X8-64G RE-MX2008-X8-64G RE-MX2008-X8-128G.

Authentication and Access Control

- **ChaCha20-Poly1305 algorithm deprecation for SSH cipher option**—The ChaCha20-Poly1305 authenticated encryption algorithm is deprecated for SSH cipher option. Configure aes-128-gcm and aes-256-gcm as the encryption algorithm for SSH Cipher option.

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

EVPN

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the `[edit protocols evpn]` hierarchy level. In most use cases, you don't need to change the default limit. If you want to change the default limit, we recommend that you don't set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

[See [mac-ip-limit](#).]

Interfaces

- Starting in Junos OS Release 24.2R1, when you run the `run show lldp local-information interface <interface-name> | display xml` command, the output is displayed under the `lldp-local-info` root tag and in the `lldp-local-interface-info` container tag. When you run the `run show lldp local-information interface | display xml` command, the `lldp-tlv-filter` and `lldp-tlv-select` information are displayed under the `lldp-local-interface-info` container tag in the output.
- **Disable keyword removal (SRX300, SRX320, SRX340, SRX345, SRX380, SRX550, SRX550M)**—The `watchdog disable` option has been removed from the `set system processes` command. You cannot configure `watchdog disable` anymore.
- **Increased limit for number of concurrent probes for real-time performance monitoring (SRX1500, SRX1600, and SRX2300, and SRX4300)**—We have increased the number of concurrent probes allowed for real-time performance monitoring (RPM) to 2000 from the previous limit of 500. [See [probe-limit](#).]

Junos OS API and Scripting

- **Changes to the XML output for ping RPCs (MX480)**—We've updated the `junos-rpc-ping` YANG module and the corresponding Junos XML RPCs to ensure that the RPC XML output conforms to the YANG schema. As a result, we changed the XML output for the following ping RPCs:
 - `<ping>`—The XML output emits `<ping-error-message>` and `<ping-warning-message>` tags instead of `<xnm:error>` and `<xnm:warning>` tags.
 - `<request-ping-ce-ip>`—The XML output is enclosed in an `<lsping-results>` root element.
 - `<request-ping-ethernet>`—
 - The `<ethping-results>` root tag includes a `<cfm-loopback-reply-entry>` or `<cfm-loopback-reply-entry-rapid>` tag for each received response. In earlier releases, a single tag enclosed all responses.
 - The XML output includes only application specific error tags and omits `<xnm:error>` tags.
 - The `<cfm-loopback-reply-entry-rapid>` tag is now reflected in the YANG schema.
 - `<request-ping-overlay>`—The `<ping-overlay-results>` element includes a new child tag `<hash-udp-src-port>`.

User Interface and Configuration

- **Configuration database maximum size increased (ACX Series, EX Series, MX Series, QFX Series, SRX Series, and vSRX)**—We've enhanced the `extend-size` statement at the `[edit system configuration-database]` hierarchy level to increase the maximum database size. On devices with a default configuration database size of ~400 MB, `extend-size` increases the maximum database size to ~2 GB. On devices with a default configuration database size of ~660 MB, `extend-size` increases the maximum database size to ~2.2 GB.

[See [configuration-database](#).]

VPNs

- **Enhancements to fix the digest option functionality for key pair generated with DSA and ECDSA (SRX Series and vSRX 3.0)**—In earlier releases, when you generated local self-signed certificates using `sha-256` digest and DSA or ECDSA encryption using `request security pki generate-key-pair certificate-id certificate-id-name size size type (dsa | ecdsa)` and `request security pki local-certificate generate-self-signed certificate-id certificate-id-name digest sha-256 domain-name domain-name subject subject-distinguished-name` commands, the generated signature always used `sha1` digest. Starting this release, the specified digest, `sha-256`, is used for the signature digest. You can verify using `show security pki local-certificate certificate-id certificate-id-name detail`
- **Enhancements to address error in generating RSA key pair with bigger key size (SRX Series)**—In earlier Junos OS releases, when you generate RSA key pair of size 4096 or greater, the command `request security pki generate-key-pair certificate-id name type rsa size 4096`, displays the error message `error: timeout communicating with pki-service daemon sometimes when PKID takes more time to respond`. Starting in Junos OS release 23.4R1, the command runs successfully without this error message.
- **Enhancements to the IKE configuration management commands in chassis cluster (SRX Series)**—In earlier Junos OS releases, in a chassis cluster mode, the following commands failed with the error message `error: IKE-Config-Management not responding to management requests on the secondary node`:
 - `show security ike statistics`
 - `show security ike sa ha-link-encryption`
 - `show security ipsec sa ha-link-encryption`
 - `show security ipsec inactive-tunnels ha-link-encryption`
 - `clear security ike sa ha-link-encryption`
 - `clear security ipsec sa ha-link-encryption`

You should run these commands only on the primary node rather than the secondary node. Starting in Junos OS Release 23.4R1, you'll not see the error message as the secondary node has no output to display.

- **Enhancements to the help string description for the threshold and interval options for VPN monitoring options (SRX Series and vSRX 3.0)**—We've enhanced the help string description of the threshold and interval options available in the configuration statement `[set security ipsec vpn-monitor-options]` to include the default values. You'll see the following description with the default values:

```
user@host# set security ipsec vpn-monitor-options ?
Possible completions:
interval Monitor interval in seconds Default :10 (2..3600 seconds)
threshold Number of consecutive failures to determine connectivity Default :10 (1..65535)
```

[See [ipsec \(Security\)](#).]

- **Enhancements to the output of show security ipsec security-associations detail command (SRX Series and vSRX 3.0)**—We've enhanced the output of `show security ipsec security-associations detail` when you enable `vpn-monitor` at the `[edit security ipsec vpn vpn-name]` hierarchy level, when your firewall runs IPsec VPN services with the new `iked` process. The output displays threshold and interval values in the command output. Starting in Junos OS Release 23.4R1, you'll notice these changes.

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

- **Enhancements to address certificate validation failures after RGO failover (SRX Series)**—Following RGO failover in the chassis cluster, you may notice that the output of the command `show services advanced-anti-malware status` displays `Requesting server certificate validation status due to CRL download failure` on the secondary node before the failover. We've made enhancements to address the issue and you'll see the following changes:
 - If there's a repeated failure to download the CRL even after multiple retry attempts, you will notice the error message `PKID_CRL_DOWNLOAD_RETRY_FAILED: CRL download for the CA failed even after multiple retry attempts, Check CRL server connection until the CRL downloads successfully`.
 - When the cluster performs a failover from the secondary to the primary node, the PKI triggers a fresh CRL download on the new primary node, resulting in successful certificate verification.
- **Reauthentication frequency recommendation for IPsec VPN with PPK (SRX Series and vSRX 3.0)**—For IPsec VPN, including the Auto Discovery VPN (ADVPN), with post-quantum pre-shared key (PPK) encryption, when the IKE security association is negotiated with the quantum keys, the `iked` process performs rekeying after 4 seconds to secure the channel. If you set the reauthentication frequency to 1, rekeying doesn't happen after 4 seconds. So we recommend you to set the reauthentication frequency to more than 1 as the first reauthentication count is used by the PPK default rekey.

[See [Quantum Safe IPsec VPN](#).]

- **Change in use of RSA signatures with SHA-1 hash algorithm**—Starting in Junos OS Release 24.2R1, there is a behavioural change by OpenSSH 8.8/8.8p1. OpenSSH 8.8/8.8p1 disables the use of RSA signatures with SHA-1 hash algorithm by default. You can use RSA signatures with SHA-256 or SHA-512 hash algorithm.

Known Limitations

IN THIS SECTION

- [Flow-Based and Packet-Based Processing](#) | 192
- [High Availability](#) | 193

Learn about known limitations in this release for SRX Series Firewalls.

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

- SRX firewalls configured in packet mode, after upgrade to Junos 24.2R1 and higher releases from earlier releases, will not forward traffic directly after the upgrade, because the packet mode configuration has changed. In the older Junos releases, when you configure **set security forwarding-options family** to use packet mode, MPLS and IPv4 (family inet) are set to packet mode automatically.

Starting in Junos OS releases 24.2R1, each address family can be configured separately to packet mode or flow mode. The default for IPv4 and IPv6 is flow mode. As a result, if your device contains the configuration `set security forwarding-options family mpls mode packet-based` before the upgrade, which sets the family IPv4 to packet mode. After the upgrade, to restore IPv4 to packet mode, you must configure `set security forwarding-options family inet mode packet-based`. Commit the configuration and reboot the device for the change to take effect.

You can check the current forwarding mode of the device using the CLI command `show security flow status`.



NOTE: No action required for IPv6 or selective packet mode using firewall filters

High Availability

- On SRX Series Firewall, Stream Control Transmission Protocol (SCTP) traffic fails in asymmetric traffic on a Multinode High Availability setup. As a workaround, you can configure the `set security forwarding-process application-services enable-sctp-port-hash hidden` command.
- On SRX5000-line firewalls do not support GTPv1 and GTPv2 for asymmetric traffic configuration. The SRX4600, SRX4300, SRX4200, SRX4100, SRX2300, SRX1600, and SRX1500 firewalls support GTPv0, GTPv1, and GTPv2 for asymmetric traffic configuration. To enable this functionality, you need to configure the `set security forwarding-process application-services enable-gtpu-distribution` setting.
- When using Advanced policy-based routing (APBR), switching routes based on the configuration is functioning correctly. However, there is an issue related to the RT_FLOW logs on the peer node in the Multinode High Availability asymmetric topology. Specifically, the logs do not include information about the uplink incoming interface name and the number of bytes/packets received on that interface.
- While the SSL proxy functionality operates smoothly on the primary node, there is an issue related to the application classification information on the peer node. We recommend you to refer the RT_FLOW session logs from the primary node to obtain accurate application details.
Example: When HTTPS traffic flows through the nodes, the active node identifies the Layer 7 standard application as IP:TCP:SSL:HTTP. However, the peer node reports the application as IP:TCP:SSL.

Open Issues

IN THIS SECTION

- [Flow-Based and Packet-Based Processing | 194](#)
- [General Routing | 194](#)
- [Platform and Infrastructure | 194](#)

Learn about open issues in this release for SRX Series Firewalls.

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

- In a chassis cluster setup configured in Active/Active mode, the fabric forward packet enters the flow module causing the flow process to pause impacting the traffic forwarding and failing the Services Processing Card (SPC). [PR1761542](#)

General Routing

- When input traffic is high and output traffic is expected equal to maximum capacity of egress interface, set the shaping explicitly equal to interface maximum capacity if default shaping does not work. [PR1712964](#)
- On SRX380 or SRX550 devices when different native VLANs are configured on the trunk interfaces between devices, you can expect packet drop. The packet drop happens because the SRX is tagging all the packets for native VLAN. [PR1750521](#)
- When peers-synchronize is configured, and master-password is configured to encrypt the configuration being synchronized. However, there is no master-password configured on the peer device, the encrypted configuration cannot be decrypted. As a workaround, configure the same master password on the peer device manually. [PR1805835](#)

Platform and Infrastructure

- On SRX1600, SRX2300, and SRX4300 Firewall devices, the MVRP registration is not working for dynamically created vlan in switching mode with MVRP protocol. [PR1804268](#)
- On SRX1600, SRX2300, and SRX4300 Firewall devices, MLD group is not formed in switching mode and multicast snooping. [PR1805291](#)
- MACsec is supported in routing mode but not in transparent mode. [PR1812427](#)
- On SRX1600, SRX2300, and SRX4300 Firewall devices, the ping from IRB to peer L3 device is not working on switching mode connecting to a L3 peer device. [PR1813712](#)

Resolved Issues

Learn about the issues fixed in this release for SRX Series Firewalls.

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

Chassis Clustering

- Core files are generated on both the nodes when you upgrade to Junos OS Release 21.2 and higher. [PR1736985](#)

Flow-Based and Packet-Based Processing

- The srpxfe process pause when ATP Cloud turned on. [PR1783101](#)
- PMI sends packets to the wrong destination. [PR1783595](#)
- Packets over GRE or IPIP or GRE(PMI) might not reach destination. [PR1791633](#)
- The GTP-U packet destination port gets duplicated to the source port and subsequently discarded by policy. [PR1798041](#)
- The commit will not go through when more than 128 VRF groups for Layer 3 VPN are configured. [PR1802089](#)
- VXLAN session not created after committing FTI configuration on both devices. [PR1807339](#)

General Routing

- TSC_DEADLINE timer feature is disabled in Intel CPUs if the microcode version is less than 0x3A. [PR1608045](#)
- High latency observed while pinging to peer device. [PR1714620](#)
- SRX4100 and SRX4200 devices accepts the datapath-debug configuration although it does not support. [PR1739559](#)
- On SRX1500, PEM alarms are displayed due to hardware limitations to read I2C. [PR1751496](#)
- ARP resolution failure for It interfaces is observed after cluster failover. [PR1753191](#)
- DNS proxy feature not working on logical tunnel interfaces. [PR1760684](#)
- Traffic drop observed right after boot up on SRX4600 device. [PR1775083](#)

- IPsec tunnel behind NAT stops passing traffic when the NAT port number or IP address changes. [PR1776216](#)
- IP monitoring fail to install route after cluster reboot. [PR1780326](#)
- Chassis alarm not present for if or var partition usage exceeds 100 percent. [PR1784983](#)
- Validate result is in processing state for more than five minutes when the configured validator port is incorrect. [PR1786432](#)
- The flowd process stops when the TLS 1.3 session ticket is received on SSL-I. [PR1788673](#)
- The srxpfe or flowd process generates core files while trying to update the path probe statistics. [PR1790782](#)
- The ISSU fails in Layer 2 HA cluster deployment. [PR1803376](#)
- The srxpfe and fwauthd process generates core files. [PR1804149](#)
- IPsec VPN is getting flapped due to warning messages [PR1805493](#)

Intrusion Detection and Prevention (IDP)

- The flowd process generates core files when the device is rebooted. [PR1786822](#)

J-Web

- J-Web default session limits have been aligned with CLI default values. [PR1788364](#)
- J-Web does not display address book entries properly after certain operations. [PR1789466](#)

Platform and Infrastructure

- Traffic loss due to PPM not offloading LACP. [PR1779749](#)
- Insufficient power alarm observed in SRX5000 Series devices. [PR1787219](#)
- The dfwd process generates core files on node1 when performing ISSU upgrade to Junos OS Release 23.1 and higher releases. [PR1794303](#)

Routing Policy and Firewall Filters

- Security policies might not synchronize during ISSU. [PR1783249](#)

Content Security

- The Web filtering does not work for HTTPS traffic that is sent from Google Chrome browser or MS Edge v124. [PR1806786](#)

User Interface and Configuration

- The SSH configuration changes do not come into affect on an existing outbound SSH client connection. [PR1791814](#)

VLAN Infrastructure

- Packet and byte counters in flow session result or traffic log are not correct for traffic uses Content Security or ALG services when SRX Series Firewall device working in Layer 2 mode. [PR1787772](#)

VPNs

- Tunnel IKE and IPsec fails to come up on SRX Series Firewalls with L2HA and FIPS after switchover. [PR1793207](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | **198**

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

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

For information about ISSU, see the [Chassis Cluster User Guide for Security Devices](#).

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

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

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



NOTE: The sixty months of support for EEOL releases is introduced in Junos OS 23.2 release and is available for all later releases. For releases prior to 23.2, the support for EEOL releases continues to be thirty six 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)	60 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](#).

Documentation Updates

This section lists the errata and changes in Junos OS Release 24.2R1 for the SRX Series documentation.

Junos OS Release Notes for vSRX

IN THIS SECTION

- [What's New | 199](#)
- [What's Changed | 204](#)
- [Known Limitations | 207](#)
- [Open Issues | 207](#)
- [Resolved Issues | 208](#)
- [Migration, Upgrade, and Downgrade Instructions | 208](#)

What's New

IN THIS SECTION

- [Application Identification \(AppID\) | 200](#)
- [High Availability | 201](#)
- [Juniper Advanced Threat Prevention Cloud \(ATP Cloud\) | 202](#)
- [Network Management and Monitoring | 202](#)
- [Platform and Infrastructure | 203](#)
- [VPNs | 203](#)

Learn about new features introduced in this release for vSRX.

Application Identification (AppID)

- **Application signature package installation enhancements (SRX Series Firewalls and vSRX)**—Starting in Junos OS Release 24.2R1, we've enhanced application signature package installation with the following changes:
 - During application signature package installation, the system performs data plane validation. This validation checks for errors in the package. If successful, the installation proceeds. If errors are found, the installation stops and reverts to the previous active version.
 - When using a chassis cluster setup, the system first installs the application signature package on the primary node and checks for any issues or problems. If the validation is successful, it then proceeds to install the same package on the secondary node.
 - The auto rollback feature now enables the system to revert to a previously working version of the application signature package. Additionally, it retains the previously designated rollback version in the event of any issues during application signature package installation.

New enhancements ensure a smooth transition by reverting to a known working version if needed.

See [[Application Signatures for Application Identification](#)].

- **CASB support (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos Release OS 24.2R1, SRX Series Firewalls support Cloud Access Security Broker (CASB).

CASB discovers SaaS applications in use and provides visibility and granular controls to protect and manage access to cloud applications. On SRX Series Firewalls, CASB provides inline activity control for the following set of cloud applications:

- Box
- Dropbox
- Salesforce
- Google Docs
- OneDrive
- SharePoint
- Slack
- Gmail

See [[Cloud Access Security Broker \(CASB\) Policy](#)].

- **SSL proxy enhancements (SRX Series Firewalls and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we introduce following enhancements for SSL proxy on SRX Series Firewalls:
 - Support of SNI extension at SSL initiation (SSL-I).
 - Support of certificate chain at SSL-I for client certificate verification.
 - Support for P-384, P-512 EC group for SSL proxy profile in addition to P-256.
 - Support for new ECDSA ciphers for SSL initiation and SSL termination profiles in non-proxy mode:
 - ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
 - ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
 - ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
 - ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
 - ECDHE_ECDSA_WITH_AES_256_CBC_SHA
 - ECDHE_ECDSA_WITH_AES_128_CBC_SHA
 - ECDHE_ECDSA_WITH_CHACHA20_POLY1305
 - New syslog messages for SSL configurations.
 - **SSL_CONFIG_MEMORY_ALLOCATION_FAILURE**— For memory allocation
 - **SSL_CONFIG_PROFILE_PROCESS_ERR** —For SSL profile processing
 - **SSL_CONFIG_CERT_PROCESS_ERR**— For SSL certificate processing.
 - **SSL_GLOBAL_CONFIG_PROCESS_ERR**—For SSL global configuration.
 - **SSL_CONFIG_PKI_IPC_ERR**— For IPC communication for SSL-PKI

See [[Cipher Suites for SSL Proxy.](#)]

High Availability

- **ADVPN support on node-local tunnels in Multinode High Availability (SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we support Auto Discovery VPN (ADVPN) on node-local tunnels configured with Multinode High Availability.

Node-local tunnels enhance Multinode High Availability by providing separate tunnels from a VPN peer device to both nodes in the setup. With ADVPN, VPN tunnels can be established dynamically between spokes. Combining ADVPN with Multinode High Availability in a node-local tunnel

deployment ensures robust network connectivity, efficient resource utilization, and seamless failover capability.

See [\[VPN Support in Multinode High Availability\]](#).

Juniper Advanced Threat Prevention Cloud (ATP Cloud)

AI-Predictive Threat Prevention leverages machine learning-based zero-day threat detection (SRX Series Firewall and vSRX Series Firewall)—Starting in Junos OS Release 24.2R1, you can configure machine learning-based threat detection for zero-day threats at line rate. File scanning during threat detection happens without Internet access and only a small section of file data is sufficient for the detection to return a verdict. Machine learning-based threat detection becomes available on your firewall when the latest antivirus signature pack is automatically downloaded from the Juniper Networks content delivery network (CDN) server to your firewall.

[See [Example: Configure Flow-Based Antivirus Policy](#), [anti-virus](#), and [show services anti-virus statistics](#).]

- **System log messages for GeoIP (SRX Series Firewalls and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we've enhanced the IP-based geolocation (GeoIP) feature to provide improved consistency checks and logging from SRX Series Firewalls that are enrolled with Juniper ATP Cloud.

The session deny message includes the following fields:

- **source-country**—Displays the country code of the source address with reference to the policy dynamic address match.
- **destination-country**—Displays the country code of the destination address with reference to the policy dynamic address match.

The system log message displays the valid country code only if the matched policy includes a dynamic address configured with GeoIP. If the matched policy does not have GeoIP configured, then the source-country and destination-country fields display N/A.

[See [System Log Explorer](#) and [Configure Juniper Advanced Threat Prevention Cloud With Geolocation IP](#).]

Network Management and Monitoring

- **Logging Infrastructure Support for RADIUS Accounting (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Starting in Junos OS Release 24.2R1, we've introduced a logging infrastructure for RADIUS accounting. The RADIUS accounting enables you to send the logs directly from dataplane to the RADIUS accounting server. When you enable RADIUS logging, logs are sent for NAT PBA ALLOC, INTERIM, and RELEASE events to the configured RADIUS server. This feature enhances the existing stream-based logging and includes:
 - Incorporation of vendor-specific attributes (VSAs) in RADIUS accounting messages

- Support multiple RADIUS accounting servers under different streams
- Manage retries and retransmissions of RADIUS accounting messages in case of failure
- Flexible and capable of supporting a backup RADIUS accounting server.

To support the feature, we've introduced the following configuration statements:

- `radius`
- `retry-count`
- `radius-accounting`
- `subscriber-extension`

Use the following commands to view and to clear the RADIUS server counters for RADIUS streams:

- `show security log radius stream`
- `clear security log radius stream.`

[See [radius \(Security Log\)](#), [retry-count \(Security Log\)](#), [radius-accounting](#), and [subscriber-extension](#).]

Platform and Infrastructure

- **Platform Upgrade (vSRX 3.0)**—Starting in Junos OS Release 24.2R1, vSRX 3.0 supports the latest upstream main FreeBSD version. This improves performance, flexibility, stability, management, and debugging.

[See [Upgrading and Downgrading to Junos with Upgraded FreeBSD](#).]

VPNs

- **Support for ChaCha20-Poly1305 algorithm (SRX1600, SRX2300, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—Starting in Junos OS Release 24.2R1, we support ChaCha20-Poly1305 authenticated encryption algorithm for IPsec VPN services. You can configure the algorithm using the option `chacha20-poly1305` for:
 - control plane with the IKEv2 protocol.
 - data plane with the IPsec ESP protocol. You configure the algorithm in PowerMode IPsec (PMI) mode for the SRX Series Firewalls, and in both the PMI and non-PMI modes for vSRX 3.0. You cannot use the algorithm for IPsec when the VPN monitoring feature is enabled.

[See [proposal \(Security IKE\)](#), [proposal \(Security IPsec\)](#), [show security ike security-associations](#), and [show security ipsec security-associations](#).]

- **Support for IPv6 address in ADVPN with ikev2 process (SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—Starting in Junos OS Release 24.2R1, we support Auto Discovery VPN (ADVPN) configuration with IPv6 address on firewalls that run the ikev2 process for IPsec VPN service.

[See [Auto Discovery VPNs](#).]

- **Support for multicast traffic in AutoVPN and ADVPN with ikev2 process (SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, and vSRX 3.0)**—Starting in Junos OS Release 24.2R1, we support IP multicast with AutoVPN and Auto Discovery VPN (ADVPN). The IP multicast uses Protocol Independent Multicast (PIM) using point-to-multipoint (P2MP) mode over st0 interface on firewalls that run the ikev2 process for IPsec VPN service. Your firewall supports IPv4 multicast in PIM sparse mode.

You can enable PIM on the st0 secure tunnel interface using the *interface-name* option at the [edit protocols pim interface *interface-name*] hierarchy level.

[See [AutoVPN](#), [Auto Discovery VPNs](#), and [interface \(Protocols PIM\)](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 204](#)
- [EVPN | 205](#)
- [User Interface and Configuration | 205](#)
- [VPNs | 205](#)

Learn about what changed in this release for vSRX.

Authentication and Access Control

- **ChaCha20-Poly1305 algorithm deprecation for SSH cipher option**—The ChaCha20-Poly1305 authenticated encryption algorithm is deprecated for SSH cipher option. Configure aes-128-gcm and aes-256-gcm as the encryption algorithm for SSH Cipher option.

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

EVPN

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the `[edit protocols evpn]` hierarchy level. In most use cases, you don't need to change the default limit. If you want to change the default limit, we recommend that you don't set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

[See [mac-ip-limit](#).]

User Interface and Configuration

- **Configuration database maximum size increased (ACX Series, EX Series, MX Series, QFX Series, SRX Series, and vSRX)**—We've enhanced the `extend-size` statement at the `[edit system configuration-database]` hierarchy level to increase the maximum database size. On devices with a default configuration database size of ~400 MB, `extend-size` increases the maximum database size to ~2 GB. On devices with a default configuration database size of ~660 MB, `extend-size` increases the maximum database size to ~2.2 GB.

[See [configuration-database](#).]

VPNs

- **Enhancements to fix the digest option functionality for key pair generated with DSA and ECDSA (SRX Series and vSRX 3.0)**—In earlier releases, when you generated local self-signed certificates using `sha-256` digest and DSA or ECDSA encryption using `request security pki generate-key-pair certificate-id certificate-id-name size size type (dsa | ecdsa)` and `request security pki local-certificate generate-self-signed certificate-id certificate-id-name digest sha-256 domain-name domain-name subject subject-distinguished-name` commands, the generated signature always used `sha1` digest. Starting this release, the specified digest, `sha-256`, is used for the signature digest. You can verify using `show security pki local-certificate certificate-id certificate-id-name detail`

- **Enhancement to the output of clear and regenerate key pair commands (vSRX 3.0)**—We've modified the output of the following commands when you clear and regenerate the same key pair to manage the secure data using hardware security module (HSM).

Starting in Junos OS 23.4R1 release, the command:

- `clear security pki key-pair certificate-id certificate-id-name` displays the message Key pair deleted successfully from the device. Key pair will be purged from the keyvault based on it's own preferences, as opposed to the message Key pair deleted successfully displayed in previous releases.
- `request security pki generate-key-pair certificate-id certificate-id-name` displays the message error:Failed to generate key pair. If the keypair was created and deleted before, please ensure that the keypair has been purged from the keyvault as opposed to the message error: Failed to generate key pair displayed in previous releases.

We made these changes to align with the cloud provider's restriction on key pair deletion, if any.

- **Enhancements to the help string description for the threshold and interval options for VPN monitoring options (SRX Series and vSRX 3.0)**—We've enhanced the help string description of the threshold and interval options available in the configuration statement `[set security ipsec vpn-monitor-options]` to include the default values. You'll see the following description with the default values:

```
user@host# set security ipsec vpn-monitor-options ?
Possible completions:
interval Monitor interval in seconds Default :10 (2..3600 seconds)
threshold Number of consecutive failures to determine connectivity Default :10 (1..65535)
```

[See [ipsec \(Security\)](#).]

- **Enhancements to the output of show security ipsec security-associations detail command (SRX Series and vSRX 3.0)**—We've enhanced the output of `show security ipsec security-associations detail` when you enable `vpn-monitor` at the `[edit security ipsec vpn vpn-name]` hierarchy level, when your firewall runs IPsec VPN services with the new `iked` process. The output displays threshold and interval values in the command output. Starting in Junos OS Release 23.4R1, you'll notice these changes.

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

- **Reauthentication frequency recommendation for IPsec VPN with PPK (SRX Series and vSRX 3.0)**—For IPsec VPN, including the Auto Discovery VPN (ADVPN), with post-quantum pre-shared key (PPK) encryption, when the IKE security association is negotiated with the quantum keys, the `iked` process performs rekeying after 4 seconds to secure the channel. If you set the reauthentication frequency to 1, rekeying doesn't happen after 4 seconds. So we recommend you to set the reauthentication frequency to more than 1 as the first reauthentication count is used by the PPK default rekey.

[See [Quantum Safe IPsec VPN](#).]

Known Limitations

There are no known limitations in hardware or software in this release for vSRX.

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

Open Issues

IN THIS SECTION

- [Flow-Based and Packet-Based Processing](#) | 207
- [VPNs](#) | 207

Learn about open issues in this release for vSRX.

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

Flow-Based and Packet-Based Processing

- In a chassis cluster setup configured in Active/Active mode, the fabric forward packet enters the flow module causing the flow process to pause impacting the traffic forwarding and failing the Services Processing Card (SPC). [PR1761542](#)

VPNs

- By default, Routing Engine in vSRX3.0 assigns only one CPU that is shared by kernel and all processes. When the CPU is very busy or hogged, it would cause some sockets in vSRX3.0 VM

timeout and got closed. As a workaround, we recommended to allocate more CPU resource to VM in HOST and Routing Engine. You can use `set security forwarding-options resource-manager cpu re to assign more CPUs to Routing Engine`. [PR1777916](#)

Resolved Issues

Learn about the issues fixed in this release for vSRX.

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

Flow-Based and Packet-Based Processing

- Traffic disruption is seen due to the segmentation violation during the DS lite tunneling. [PR1779792](#)

VPNs

- Encapsulation or de-encapsulation might not work correctly when PMI is enabled and using life-sizes. [PR1758785](#)
- IPsec tunnel behind NAT stops passing traffic when the NAT port number or IP address changes. [PR1776216](#)
- The ike idle-time tears down VPN even when active traffic pass through tunnel. [PR1802145](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 215

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.

For SRX Series Firewall and vSRX Junos OS upgrade path information, see [Junos upgrade paths for SRX platforms](#).

You also can upgrade to Junos OS Release 24.2R1 for vSRX using J-Web (see [J-Web](#)) or the Junos Space Network Management Platform (see [Junos Space](#)).

The following limitations apply:

- 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 24.2R1 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
24.2K	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/eedebg_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 24.2R1 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-x86-64-24.2-2024-06-06.0_RELEASE_24.2_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-
x86-64-24.2-2024-06-06.0_RELEASE_24.2_THROTTLE.tgz no-copy no-validate reboot
Verified junos-vsrx-x86-64-24.2-2024-06-06.0_RELEASE_24.2_THROTTLE signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING: This package will load JUNOS 24.2 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-24.2-2024-06-06.0_RELEASE_24.2_THROTTLE.tgz

```

```

Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31 junos-srx-mr-
vsrx-24.2-2024-06-06.0_RELEASE_24.2_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31 junos-srx-mr-
vsrx-24.2-2024-06-06.0_RELEASE_24.2_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-24.2-2024-06-06.0_RELEASE_24.2_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-24.2-2024-06-06.0_RELEASE_24.2_THROTTLE-linux.tgz ...
upgrade_platform: Input package /var/tmp/junos-srx-mr-
vsrx-24.2-2024-06-06.0_RELEASE_24.2_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-24.2-2024-06-06.0_RELEASE_24.2_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-24.2-2024-06-06.0_RELEASE_24.2_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 24.2R1 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 24.2-2024-06-06.0_RELEASE_24.2_THROTTLE Kernel 64-bit
JNPR-11.0-20240606.170745_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: 24.2-2024-06-06.0_RELEASE_24.2_THROTTLE
JUNOS OS Kernel 64-bit [20240606.170745_fbsd-builder_stable_11]
JUNOS OS libs [20240606.170745_fbsd-builder_stable_11]
JUNOS OS runtime [20240606.170745_fbsd-builder_stable_11]
JUNOS OS time zone information [20240606.170745_fbsd-builder_stable_11]
JUNOS OS libs compat32 [20240606.170745_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20240606.170745_fbsd-builder_stable_11]
JUNOS py extensions [20240606.110007_ssd-builder_release_174_throttle]
JUNOS py base [20240606.110007_ssd-builder_release_174_throttle]
JUNOS OS vmguest [20240606.170745_fbsd-builder_stable_11]
JUNOS OS crypto [20240606.170745_fbsd-builder_stable_11]
JUNOS network stack and utilities [20240606.110007_ssd-builder_release_174_throttle]
JUNOS libs [20240606.110007_ssd-builder_release_174_throttle]
JUNOS libs compat32 [20240606.110007_ssd-builder_release_174_throttle]
JUNOS runtime [20240606.110007_ssd-builder_release_174_throttle]
JUNOS Web Management Platform Package [20240606.110007_ssd-builder_release_174_throttle]
JUNOS srx libs compat32 [20240606.110007_ssd-builder_release_174_throttle]
JUNOS srx runtime [20240606.110007_ssd-builder_release_174_throttle]
JUNOS common platform support [20240606.110007_ssd-builder_release_174_throttle]
JUNOS srx platform support [20240606.110007_ssd-builder_release_174_throttle]
JUNOS mtx network modules [20240606.110007_ssd-builder_release_174_throttle]
JUNOS modules [20240606.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp modules [20240606.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp libs [20240606.110007_ssd-builder_release_174_throttle]
JUNOS srx libs [20240606.110007_ssd-builder_release_174_throttle]
JUNOS srx Data Plane Crypto Support [20240606.110007_ssd-builder_release_174_throttle]
JUNOS daemons [20240606.110007_ssd-builder_release_174_throttle]
JUNOS srx daemons [20240606.110007_ssd-builder_release_174_throttle]
JUNOS Online Documentation [20240606.110007_ssd-builder_release_174_throttle]
JUNOS jail runtime [20240606.170745_fbsd-builder_stable_11]
JUNOS FIPS mode utilities [20240606.110007_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 sixty months after the first general availability date and customer support for an additional six more months.



NOTE: The sixty months of support for EEOL releases is introduced in Junos OS 23.2 release and is available for all later releases. For releases prior to 23.2, the support for EEOL releases continues to be thirty six 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 14: 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)	60 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 | 217
- Creating a Service Request with JTAC | 218

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- **JTAC policies**—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- **Product warranties**—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- **JTAC hours of operation**—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>

- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>

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.

- 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

02 July 2024—Revision 1, MX-only Junos OS Release 24.2R1.

04 July 2024—Revision 2, cRPD and MX only Junos OS Release 24.2R1.

19 July 2024—Revision 3, Junos OS Release 24.2R1.

26 July 2024—Revision 4, Junos OS Release 24.2R1.

02 August 2024—Revision 5, Junos OS Release 24.2R1.

08 August 2024—Revision 6, Junos OS Release 24.2R1.

15 August 2024—Revision 7, Junos OS Release 24.2R1.

27 September 2024—Revision 8, Junos OS Release 24.2R1.

05 October 2024—Revision 9, Junos OS Release 24.2R1.

23 October 2024—Revision 10, Junos OS Release 24.2R1.

07 November 2024—Revision 11, Junos OS Release 24.2R1.

13 December 2024—Revision 12, Junos OS Release 24.2R1.

20 December 2024—Revision 13, Junos OS Release 24.2R1.

17 January 2025—Revision 14, Junos OS Release 24.2R1.

04 February 2025—Revision 15, Junos OS Release 24.2R1.

27 February 2025—Revision 16, Junos OS Release 24.2R1.

31 March 2025—Revision 17, Junos OS Release 24.2R1.

18 September 2025—Revision 18, Junos OS Release 24.2R1.

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 © 2025 Juniper Networks, Inc. All rights reserved.