

Release Notes

Published
2026-09-24

Junos OS Release 25.4R1®

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, MX Series, NFX Series, QFX Series, SRX Series, and vSRX. These release notes accompany Junos OS Release 25.4R1. 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

Junos OS Release Notes for ACX Series

What's New | 1

- Junos Telemetry | 1
- Post-Quantum Cryptography (PQC) | 2
- Routing Protocols | 2
- Additional Features | 2

What's Changed | 3

Known Limitations | 7

Open Issues | 7

Resolved Issues | 7

Migration, Upgrade, and Downgrade Instructions | 8

- Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 9

Junos OS Release Notes for cRPD

What's New | 10

- Juniper Extension Toolkit (JET) | 10

What's Changed | 11

Known Limitations | 11

Open Issues | 11

Resolved Issues | 12

Junos OS Release Notes for cSRX

What's New | 12

- Application Layer Gateways (ALGs) | 13
- Device Security | 13

Junos OS API and Scripting | 13

Platform and Infrastructure | 14

What's Changed | 14

Known Limitations | 15

Open Issues | 15

Resolved Issues | 15

Junos OS Release Notes for EX Series

What's New in 25.4R1-S1 | 16

Hardware | 16

Chassis | 22

Interfaces | 22

What's New in 25.4R1 | 22

EVPN | 23

Interfaces | 24

Junos OS API and Scripting | 25

Junos Telemetry | 25

MAC Learning | 27

Network Management and Monitoring | 27

Post-Quantum Cryptography (PQC) | 28

Software Installation and Upgrade | 29

Additional Features | 29

What's Changed | 29

Known Limitations | 35

Open Issues | 36

Resolved Issues | 38

Migration, Upgrade, and Downgrade Instructions | 41

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 41

Junos OS Release Notes for JRR Series

What's New | 42

What's Changed | 43

Known Limitations | 43

Open Issues | 43

Resolved Issues | 44

Migration, Upgrade, and Downgrade Instructions | 44

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 44

Junos OS Release Notes for Juniper Secure Connect

What's New | 46

What's Changed | 46

Known Limitations | 46

Open Issues | 46

Resolved Issues | 47

Junos OS Release Notes for MX Series

What's New in 25.4R1-S1 | 47

Hardware | 48

MACsec | 57

What's New in 25.4R1 | 58

Hardware | 58

Chassis | 68

EVPN | 68

High Availability | 69

Interfaces | 70

Junos OS API and Scripting | 70

Precision Time Protocol (PTP) | 71

Junos Telemetry | 73

MACsec | 75

MPLS | 75

Network Management and Monitoring | 76

NextGen Port Extender (NGPE) | 77

Post-Quantum Cryptography (PQC) | 77

Routing Protocols | 77

Services Applications | 79

Source Packet Routing in Networking (SPRING) or Segment Routing | 80

Subscriber Management and Services | 82

Additional Features | 83

What's Changed | 84

Known Limitations | 91

Open Issues | 92

Resolved Issues | 99

Migration, Upgrade, and Downgrade Instructions | 108

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 112

Junos OS Release Notes for NFX Series

What's New | 114

Application Layer Gateways (ALGs) | 114

What's Changed | 114

Known Limitations | 115

Open Issues | 115

Resolved Issues | 116

Migration, Upgrade, and Downgrade Instructions | 117

Junos OS Release Notes for QFX Series

What's New | 120

Class of Service | 120

EVPN | 120

Flow-Based and Packet-Based Processing | 122

Junos OS API and Scripting | 123

Junos Telemetry | 123

MAC Learning | 124

Multicast | 124

Network Address Translation (NAT) | 124

Post-Quantum Cryptography (PQC) | 125

Routing Policy and Firewall Filters | 125

Routing Protocols | 125

Additional Features | 127

What's Changed | 127

Known Limitations | 133

Open Issues | 133

Resolved Issues | 134

Migration, Upgrade, and Downgrade Instructions | 136

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life
Releases | 148

Junos OS Release Notes for SRX Series

What's New | 149

Application Identification (AppID) | 150

Application Layer Gateways (ALGs)	150
Chassis	150
Device Security	151
High Availability	151
Identity Aware Firewall	152
Interfaces	153
Juniper Advanced Threat Prevention Cloud (ATP Cloud)	153
Junos OS API and Scripting	153
Junos Traffic Vision	154
MACsec	155
MPLS	155
Network Management and Monitoring	156
Post-Quantum Cryptography (PQC)	156
Public Key Infrastructure (PKI)	157
Additional Features	157

What's Changed | 158

Known Limitations | 165

Open Issues | 166

Resolved Issues | 166

Migration, Upgrade, and Downgrade Instructions | 170

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases	170
--	-----

Junos OS Release Notes for vSRX

What's New | 172

Application Identification (AppID)	172
Application Layer Gateways (ALGs)	172

Device Security | 173

High Availability | 173

Identity Aware Firewall | 174

Junos OS API and Scripting | 174

Junos Traffic Vision | 175

Public Key Infrastructure (PKI) | 176

What's Changed | 176

Known Limitations | 179

Open Issues | 179

Resolved Issues | 179

Migration, Upgrade, and Downgrade Instructions | 180

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life
Releases | 186

Licensing | 187

Finding More Information | 188

Requesting Technical Support | 189

Revision History | 190

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [What's New | 1](#)
- [What's Changed | 3](#)
- [Known Limitations | 7](#)
- [Open Issues | 7](#)
- [Resolved Issues | 7](#)
- [Migration, Upgrade, and Downgrade Instructions | 8](#)

What's New

IN THIS SECTION

- [Junos Telemetry | 1](#)
- [Post-Quantum Cryptography \(PQC\) | 2](#)
- [Routing Protocols | 2](#)
- [Additional Features | 2](#)

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

Junos Telemetry

- **Sensor path support for the table connection `disable-metric-propagation` leaf (ACX710, MX10008, and MX10016)**—By default, the OpenConfig protocol sets the destination protocol metric based on the source protocol metric. Set the `disable-metric-propagation` leaf to true to stop this behavior. The device then sets the metric to zero or a policy-defined value. Subscribe to the following resource paths to obtain sensor-specific information:
 - `/network-instances/network-instance/table-connections/table-connection/config/disable-metric-propagation`

- `/network-instances/network-instance/table-connections/table-connection/state/disable-metric-propagation`

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

Post-Quantum Cryptography (PQC)

- **PQC signatures for software images with off-box verification (ACX Series, EX Series, MX Series, QFX Series, and SRX Series)**—Use post-quantum cryptography (PQC) signatures to ensure software images remain unaltered, protecting integrity and guarding against quantum threats. The images comply with algorithms recommended by Commercial National Security Algorithm 2.0 (CNSA 2.0):
 - ML-DSA-87 PQC algorithm for digital signatures
 - SHA-512 for hashing

For off-box verification, request the PQC signature from the support portal. Confirm the image's SHA-512 hash, retrieve the public key from the certificate, and validate the signature with your chosen verifier. PQC signatures provide additional security beyond existing legacy signatures.

[See [PQC Signatures for Software Images](#).]

Routing Protocols

- **IS-IS multi-instance support on a single interface (ACX5448, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX10002-60C, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016)**—We have enhanced the IS-IS multi-instance feature to support multiple IS-IS instances on the same logical interface with instance identifier TLV 7.

Include the `instance-id` statement at the `[edit protocols isis-instance name hierarchy level`.

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

Additional Features

We've extended support for the following features to these platforms.

- **Renaming OpenSSH implementation to JSSH (all platforms)**—The OpenSSH implementation in Junos OS is renamed "JSSH" to highlight its customized nature.

[See [ssh](#).]

What's Changed

IN THIS SECTION

- [Dynamic Host Configuration Protocol | 3](#)
- [General Routing | 3](#)
- [MPLS | 5](#)
- [Network Management and Monitoring | 5](#)
- [Platform and Infrastructure | 5](#)
- [Routing Protocols | 5](#)
- [Subscriber Access Management | 6](#)
- [User Interface and Configuration | 6](#)

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

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

General Routing

- **Control Maximum 802.1X Client Connections per Interface**—By default, dot1x interfaces configured in multiple supplicant mode have a client limit of 100 authenticated connections per interface. Any additional connection attempts beyond this limit will be automatically blocked.
- **New option for debug collector data storage path**—We've included the `outdir` option to specify an output directory for storing debug collector data in a customised path. This allows you to organise and access diagnostic information more efficiently, adapting storage to your specific requirements.
- A new counter `Sessions hit due to high rate` is added to the `show services service-sets screen-session-limit-counters` command for all subscriber traffic. This counter tracks the sessions that come up on the screen irrespective of the `alarm-without-drop` configuration. When the `alarm-without-drop` option is

disabled, all the counters display updated statistics. When `alarm-without-drop` is enabled, then, the screen-drop counters on the `show services service-sets statistic screen-drop` command do not increase. The `Sessions hit due to high rate` value is displayed.

When you run the `request vmhost zeroize` command to zeroize a single Routing Engine on a dual Routing Engine device, the CLI incorrectly displays a message indicating that it will zeroize both Routing Engines

- **Validation of `/vm-primary` mount during JDM installation or upgrade (Junos Node Slicing External Server Deployment)**—During installation or upgrade of the Juniper Device Manager (JDM) on external servers running Junos Node Slicing Release 25.2 or later, an issue occurred with the validation of the `/vm-primary` mount that stores the GNF images. When `/vm-primary` was mounted using logical volumes (LVM), JDM would fail to detect that the underlying storage was an SSD. This issue is now fixed. However, the fix introduces a new dependency on the LVM2 package in the host OS. This package is included by default in standard installations of both RHEL and Ubuntu external servers. However, it is advised that you check if the LVM2 package is already installed on the host before installing or upgrading JDM.
- **Chassis redundancy failover options are not supported on ACX7348, ACX7332, and ACX7059 platforms**—The `disk-read-threshold`, `disk-write-threshold`, `not-on-disk-underperform`, `on-disk-failure`, and `on-loss-of-keepalives` statements under chassis redundancy failover are not supported on ACX7348, ACX7332, and ACX7059 platforms. Configuration attempts using these statements are rejected. This change is platform-specific and does not affect other Junos platforms where these statements may continue to be supported.[PR1885312](#)
- **Static flowspec prefix validation**—Commit validation now rejects static flowspec destination and source prefixes that use a shorter mask than the configured address family permits. The configuration must use a valid prefix length before the commit can succeed. This alignment with static route validation prevents invalid flowspec entries from being installed and ensures consistent routing policy enforcement across IPv4 and IPv6 configurations.[PR1903233](#)
- **Enhanced ping command XML output (Junos OS and Junos OS Evolved)**—The XML output for the CE-IP ping has enhanced details to differentiate between successful and failed ping attempts. The update introduces new tags that track the total number of packets sent and received, while also highlighting host unreachable replies and the overall packet loss percentage.

[See [Supported and Unsupported Features for CE-IP Ping](#).][PR1906967](#)

- **`request jnu role delete` Command Unhidden (Junos OS and Junos OS Evolved)**—The `request jnu role delete` command, which was previously hidden, is now available for use. This command allows users to delete the role of controllers or satellites directly through the CLI. This enhancement improves administrative control and flexibility in role management within the network infrastructure.

[See [request jnu role](#).][PR1930291](#)

MPLS

- **Label field naming update in `show route forwarding-table` extensive output (Junos OS and Junos OS Evolved)**—The CLI output for MPLS forwarding entries now uses "Label" instead of "Load balance label." This name change standardizes the field name and improves overall clarity when reviewing MPLS forwarding information. The change affects only the output text and does not modify label allocation, forwarding behavior, or traffic handling. No configuration changes are required.

Network Management and Monitoring

- **Deprecation of `shell` option**—The `shell` option no longer requires a separate configuration and is now the default behavior. Deprecating the `shell` option enhances efficiency and simplifies management tasks.
- **Shell Command Logging Enhancement**—All shell commands executed on the device and root sessions are now logged by default. This enhancement ensures enhanced security and auditability by capturing all commands entered in any shell environment, preventing bypassing audit logging. [PR1867216](#)

Platform and Infrastructure

- **Commit validation for unique user IDs**—We have added support to validate the user configuration to ensure that each user is assigned a unique UID. A commit fails if duplicate UIDs are detected, ensuring stronger validation and preventing identity conflicts. Previously, a commit was successful even when multiple users shared the same UID, triggering only a warning and logging a syslog message. [PR1890656](#)

Routing Protocols

- **SRv6 MSD advertisement compliance**—MSD type 43 (SRH Max T.insert) is no longer advertised in IS-IS SRv6 node capability TLVs to match the latest standards draft. This alignment maintains interoperability with peer devices and tools that expect the reduced MSD type list and ensures SRH limits are reported consistently across the network. [PR1879051](#)
- **Auto BGP peer discovery for unnumbered interfaces of iBGP peers**—Junos supports single-hop iBGP sessions for unnumbered peering. Previously, the iBGP implementation created an indirect next-hop for route resolution, even if the iBGP peer was directly connected. [PR1884861](#)
- **BGP route-refresh warning messages added**—New messages have been introduced to the `show bgp diagnostics warnings` command output. These messages provide diagnostic information when the number of BGP route-refresh messages exceed predetermined thresholds within specific time intervals of 5 minutes, 30 minutes, and 1 hour.

[See [show bgp diagnostics warnings](#).] [PR1909467](#)

- **BGP UPDATE messages must have the marker field set to all ones**—The marker field in BGP header is a 16-octet field included for compatibility that must be set to all ones. All Junos OS and Junos OS Evolved devices send a notification and tear down the BGP session if the marker field in the header of the received update message is not set to all ones. [PR1955678](#)

Subscriber Access Management

- **Addition of message-authenticator and no-message-authenticator attributes under access radius-server, access profile radius-server, and system radius-server hierarchies**—Set message-authenticator if you require the RADIUS server to include the Message Authenticator attribute in replies to Access-Request messages. Set no-message-authenticator to not require that attribute. [PR1871147](#)

User Interface and Configuration

- **Stale ui-state.db data in persistent NETCONF sessions post-mgd restart**—Existing NETCONF sessions might fetch stale data from ui-state.db after mgd -N restart. New sessions correctly map the refreshed database. Scripts must establish new sessions post-restart to access updated values. Functional configuration remains unaffected. Script failures monitoring "local-host" NETCONF sessions—Scripts might fail when including "local-host" NETCONF sessions in monitoring operations. Internal sessions are now excluded from tracking. Scripts must filter out "local-host" sessions. No impact to internal application functionality.
- **Generate genstate YANG modules on Junos devices**—You can use show system schema operational command or equivalent RPC to generate the genstate YANG modules in the specified output directory on a device.

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

- **Updated Annotate Command Behavior**—The annotate command now correctly handles multi-line comments that start with `#`. This change ensures that all lines of a multi-line comment remain comments, preventing unauthorized configuration changes when the configuration is committed, reloaded, or rolled back. These improvements enhance security by mitigating potential privilege escalation attacks from users with limited configuration editing permissions. [PR1868636](#)
- **Enhanced Permission Checks for Rename/Copy Operations**—New permission checks have been introduced for rename and copy operations within the configuration hierarchy. These checks parse the hierarchy being modified to ensure the user has the required permissions for the hierarchy and its sub-hierarchies. If permissions are insufficient, the operation will fail and a "Permission denied" error will be displayed. This enhancement ensures that configuration modifications are performed only by authorized users, improving security and preventing unintended changes to critical system settings. [PR1882303](#)

Known Limitations

IN THIS SECTION

- [Infrastructure | 7](#)

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.

Infrastructure

- When you upgrade from Junos OS Release 21.2 release and prior to Junos OS Release 21.2 and onward, validation and upgrade might fail. You must use the `no-validate` option to complete the upgrade successfully. [PR1568757](#)

Open Issues

There are no known issues in hardware or software 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.

Resolved Issues

IN THIS SECTION

- [General Routing | 8](#)

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

- Packet loss occurs when explicit Null is disabled for BGP-LU routes in ECMP scenarios. [PR1881742](#)
- EVPN-MPLS BUM Traffic Disruption occurs due to Incorrect QinQ STag Insertion. [PR1882561](#)
- Optics fails to come up post reboot. [PR1887528](#)
- SFP-T port does not come up after system restart. [PR1896458](#)
- Traffic loss occurs with multiple VRRP instances when you use the same VRRP group ID and one of VRRP instance transitions backup. [PR1897117](#)
- ARP resolution and device discovery failure occurs due to unexpected VLAN tags on ARP replies. [PR1897336](#)
- Deactivating or deleting the auto-configure command in an interface does not work properly. [PR1902855](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 9

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

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

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

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

- Standard End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for 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 1: EOL and EEOL Releases

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

IN THIS SECTION

- [What's New | 10](#)
- [What's Changed | 11](#)
- [Known Limitations | 11](#)
- [Open Issues | 11](#)
- [Resolved Issues | 12](#)

What's New

IN THIS SECTION

- [Juniper Extension Toolkit \(JET\) | 10](#)

Learn about new features introduced in this release for cRPD.

Juniper Extension Toolkit (JET)

- **Enhanced PRPD APIs to support remotely programmed EVPN Type 5 routes (cRPD)**—Juniper Extension Toolkit (JET) APIs allow you to remotely program static BGP routes and direct traffic through your network using an external controller such as a JET client. You can remotely program EVPN Type 5 routes using all existing BGP APIs. These APIs support EVPN Type 5 addresses and BGP communities. To enable the router MAC community, use the format `router-mac:mac-address`. To enable the encapsulation community with a VXLAN, use the format `encapsulation:vxlan`.

To get improved notifications, use the `RouteRibLocalSubscribe` API in the BGP Route Service RPC. This RPC streams all BGP communities, including extended BGP communities, to the local routing table in the controller. You can specify a policy for each route subscription request. These policies can match a route family, route distinguisher, community, or prefix. You can find details in version 2 of the `jnx_routing_bgp_service.proto` file.

[See [Overview of JET APIs](#) and [Controller-Based BGP Multicast Signaling](#).]

What's Changed

IN THIS SECTION

- [Dynamic Host Configuration Protocol](#) | 11

Learn about what changed in this release for cRPD.

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

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 open 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 | 12](#)
- [What's Changed | 14](#)
- [Known Limitations | 15](#)
- [Open Issues | 15](#)
- [Resolved Issues | 15](#)

What's New

IN THIS SECTION

- [Application Layer Gateways \(ALGs\) | 13](#)
- [Device Security | 13](#)
- [Junos OS API and Scripting | 13](#)
- [Platform and Infrastructure | 14](#)

Learn about new features introduced in this release for cSRX.

Application Layer Gateways (ALGs)

- **Support for client identifier in forwarded DNS queries using experimental EDNS(0) option (NFX150, NFX250, NFX350, cSRX, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—You can enhance DNS query responses by utilizing the *Client ID in Forwarded DNS Queries* feature on SRX Series devices. This functionality supports precise client identification, essential for services such as parental control. You can configure SRX Series Firewalls to include client identifiers such as MAC addresses, IPv4, or IPv6 addresses in DNS queries. This feature, intended for controlled environments, generates targeted DNS responses based on the originating device's identity, improving service accuracy and network efficiency.

[See [DNS ALG](#).]

Device Security

- **FQDN ID for enhanced policy management and dynamic IP resolution (SRX Series, cSRX, and vSRX)**—Use unique fully qualified domain name (FQDN) ID mappings to manage frequent IP address changes. The system stores each FQDN's identifier (ID) in the Routing Engine and Packet Forwarding Engine. This FQDN ID storage enables quick lookups without constant policy updates, improving stability. This feature runs by default.

[See [DNS Snooping for Security Policies](#).]

Junos OS API and Scripting

- **Support for libslax 3.1.6 and SLAX version 1.3 (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—We've upgraded the libslax library to version 3.1.6, which corresponds to SLAX version 1.3. The upgraded library includes:
 - Slax processor enhancements including a new mode, additional options, and simplified argument parsing

- New libslax extension library functions
- Improved SLAX syntax options
- New SLAX functions and enhancements to existing functions and statements
- Hexadecimal numbers support in SLAX scripts

This update aligns SLAX processing with contemporary scripting conventions, ensuring efficient, readable scripting while maintaining backward compatibility.

[See [libslax Distribution Overview](#).]

Platform and Infrastructure

- **cSRX deployment on Kubernetes using CRI-O or Podman instead of Docker**—You can deploy cSRX in Kubernetes using Container Runtime Interface-Open (CRI-O) or Podman as the container runtime instead of Docker. This runtime choice maintains compatibility and operational continuity as Kubernetes deprecates Docker support. Configure your Kubernetes cluster to use CRI-O or Podman and deploy cSRX with those runtimes.

[See [cSRX Container Firewall with Kubernetes](#).]

- **Integration of Mellanox ConnectX SmartNIC in containerized deployments (cSRX)**—You can deploy cSRX with single-root I/O virtualization (SR-IOV) on SmartNICs using Data Plane Development Kit (DPDK) poll mode drivers on x86 and ARM platforms to enhance network performance. This approach also facilitates the transition from virtualized network function (VNF) to cloud-native network function (CNF) infrastructure.

This integration, supported by the DPDK driver modules, addresses the increasing demand for Mellanox NICs, enabling you to meet the evolving requirements of your infrastructure efficiently.

[See [Requirements for cSRX Container Firewall](#).]

What's Changed

IN THIS SECTION

- [Dynamic Host Configuration Protocol | 15](#)

Learn about what changed in this release for cSRX.

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

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 in 25.4R1-S1 | 16](#)
- [What's New in 25.4R1 | 22](#)
- [What's Changed | 29](#)
- [Known Limitations | 35](#)
- [Open Issues | 36](#)
- [Resolved Issues | 38](#)
- [Migration, Upgrade, and Downgrade Instructions | 41](#)

What's New in 25.4R1-S1

IN THIS SECTION

- [Hardware | 16](#)
- [Chassis | 22](#)
- [Interfaces | 22](#)

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

Hardware

- **New EX4100-H-12T Switch (EX Series)** - A new 12-port, ruggedized, non-PoE 3.5 RU switch in the EX4100-H series, designed with support for a wide range DC input voltage. The switch offers 12 x 1G Ethernet ports, features a fanless design, and supports an extended operating temperature range of -40°C through 75°C, suitable for various industrial conditions.

Table 2: Features Supported on EX4100-H-12T Switches

Feature	Description
Access and authentication	• Support for 802.1X authentication. [See 802.1X Authentication.] • Support for captive portal authentication. [See Captive Portal Authentication.]
Class of Service (CoS)	• Support for CoS configuration. [See Junos OS CoS for EX Series Switches Overview.]
DHCP	Supports the following DHCP features: • DHCPv4 and DHCPv6 client • DHCPv4 and DHCPv6 server • DHCPv4 and DHCPv6 relay agent [See DHCP User Guide.]

Table 2: Features Supported on EX4100-H-12T Switches (*Continued*)

Feature	Description
EVPN	• Support for EVPN-VXLAN group-based policies. [See Micro and Macro Segmentation using Group Based Policy in a VXLAN.] • Supports Layer 2 VXLAN gateway services in an EVPN-VXLAN network. [See EVPN User Guide.] • Support for Layer 3 VXLAN gateway in EVPN-VXLAN centrally routed bridging (CRB) overlay or edge-routed bridging (ERB) overlay networks on standalone switches or Virtual Chassis. [See EVPN User Guide.] The Virtual Chassis doesn't support EVPN-VXLAN multihoming. You can use the standalone switch as an EVPN-VXLAN provider edge (PE) device in multihoming use cases. We support the following Layer 2 VXLAN gateway features in an EVPN-VXLAN network: • Active/active multihoming • Proxy ARP use and ARP suppression, and Neighbor Discovery Protocol (NDP) use and NDP suppression on non-IRB interfaces • Ingress node replication for broadcast, unknown unicast, and multicast (BUM) traffic forwarding [See EVPN User Guide.]
Hardware	• The EX4100-H-12T has the following port configuration: • Twelve 1 Gigabit-Ethernet (GbE) non-PoE BASE-T ports • Two 1/10 Gbps SFP/SFP+ MACsec-enabled uplink ports • Two 1/10 Gbps SFP/SFP+ VCPs. MACsec supported in Network mode. [See EX4100-H-12T Hardware Guide.]

Table 2: Features Supported on EX4100-H-12T Switches (*Continued*)

Feature	Description
High availability and resiliency	Resiliency support for inter-integrated circuit (I2C), disk failure, and disk health. [See High Availability User Guide.]
Layer 2 features	- Support for Layer 2 features. [See Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation, Layer 2 Bridge Domains Overview, and Understanding Layer 2 Learning and Forwarding.] - Support for Layer 2 multicast features. [See Multicast Overview and Understanding Multicast Snooping.]
Layer 3 features	Support for Layer 3 features and interior gateway protocols (OSPF, IS-IS, RIP, and ECMP) for IPv4 and IPv6. [See Understanding OSPF Configurations and BGP Overview.]
MACsec	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. [See Configuring MACsec.]

Table 2: Features Supported on EX4100-H-12T Switches (*Continued*)

Feature	Description
Network management and monitoring	• Support for Ethernet OAM link fault management (LFM) and CFM features. [See Ethernet OAM and CFM for Switches and OAM Link Fault Management.] • Support for IEEE 802.1ag CFM on service provider interfaces and Q-in-Q (point-to-point) interfaces. [See Introduction to OAM Connectivity Fault Management (CFM).] • Support for Juniper Mist Wired Assurance. [See Juniper AI-Driven Enterprise and Overview of EX Series Switches and the Juniper Mist Cloud.] • Support for: • Spanning-tree protocols. [See Spanning Tree Protocol Instances and Interfaces.] • sFlow network monitoring technology. [See sFlow Monitoring Technology.] • Local and remote port mirroring, and remote port mirroring to an IP address (GRE encapsulation). [See Port Mirroring and Analyzers.]

Table 2: Features Supported on EX4100-H-12T Switches (*Continued*)

Feature	Description
Software installation and upgrade	• Support for the PHC. Obtaining Configurations and Software Image Without User Intervention Using Phone-Home Client • Secure boot and common BIOS support (EX4100-H-12T). The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms. [See Junos OS Overview.] • Firmware upgrade support. (EX4100-H-12T) [See request system firmware upgrade.] • Support for ZTP. [See Zero Touch Provisioning.] • Support for SZTP. [See Secure Zero Touch Provisioning and Generate Voucher Certificate.]
Timing	• Support for Precision Time Protocol (PTP) transparent clock. [See PTP Transparent Clocks.]
Uplink failure detection	• Support for debounce interval configuration. You can configure the debounce interval, which is the time (in seconds) that elapses before the downlink interfaces are brought up after a state change of the uplink interfaces. [See Uplink Failure Detection.]
Virtual Chassis	• Support for Virtual Chassis. [See Understanding EX Series Virtual Chassis.]

Chassis

- **Supporting Chassis, FRU, Environment Monitoring, and Systems Alarms for EX series switch (EX4100-H-12T)**—The EX4100-H-12T switch monitors PEM modules, board temperature sensors. It has a variable DC input voltage range of 9.6VDC to 60VDC. The power is provided with external PSUs or power source. N+1 redundancy is supported. The switch also has the capability to display input power supplied to each PSU input slot. Watchdog timer is enabled in the h/w and in the core OS and the watchdog patting logic is implemented for periodically patting the watchdog.

Interfaces

- **Interfaces Support (EX4100-H-12T)**—You can deploy EX4100-H-12T with flexible port options across three PICs. Use 12 downlink ports (PIC 0) at 10/100/1000 Mbps for access connectivity. Use two stacking ports (PIC 1) in network mode at 100 Mbps, 1 Gbps, or 10 Gbps. Select Virtual Chassis (HGoE) mode at 10 Gbps to build stackable architectures. Use two uplink ports (PIC 2) in network mode at 100 Mbps, 1 Gbps, or 10 Gbps. Select VC (HGoE) mode at 10 Gbps to scale uplinks or interconnects. See [Port Speed on EX Series Switches](#).

What's New in 25.4R1

IN THIS SECTION

- [EVPN | 23](#)
- [Interfaces | 24](#)
- [Junos OS API and Scripting | 25](#)
- [Junos Telemetry | 25](#)
- [MAC Learning | 27](#)
- [Network Management and Monitoring | 27](#)
- [Post-Quantum Cryptography \(PQC\) | 28](#)
- [Software Installation and Upgrade | 29](#)
- [Additional Features | 29](#)

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



NOTE: The EX2300 and EX3400 models are documented but not supported in this release.

EVPN

- **Unified access policy (EX4100, EX4400, EX4650, and QFX5120)**—Unified access policy extends group-based policy (GBP) support to Mist APs, including to parts of the wired and wireless access network outside of the EVPN-VXLAN infrastructure. GBP tags are learned through proprietary control plane messages from Mist APs and across access switches, allowing both wired and wireless clients to participate in GBP microsegmentation.

[See [Microsegmentation Using Group-Based Policies](#).]

- **GBP support for DHCP, ARP, and neighbor discovery packets when snooping and inspection are enabled (EX4100, EX4400, EX4650, and QFX5120)**—DHCP snooping, dynamic ARP inspection, and dynamic IPv6 neighbor discovery inspection now include GBP support for DHCP, Address Resolution Protocol (ARP), and neighbor discovery packets, respectively. Previously, when snooping and inspection were enabled, GBP processing of the snooped and inspected packets did not take place.

[See [Microsegmentation Using Group-Based Policies](#).]

- **GBP on an IPv6 underlay (EX4100, EX4400, EX4650, and QFX5120)**—Group-based policy (GBP) is now supported on top of an IPv6 underlay network. With an IPv6 underlay, you can take advantage of the expanded addressing capabilities and efficient packet processing that the IPv6 protocol offers.

[See [Microsegmentation Using Group-Based Policies](#) and [EVPN-VXLAN with an IPv6 Underlay](#).]

- **EVPN maintenance mode CLI for multihomed ERB leaf nodes (EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200, and QFX5210)**—You can streamline the upgrade process for EVPN-VXLAN leaf devices by utilizing the *maintenance mode* CLI. This feature enables you to isolate multihomed nodes and manage the upgrade with minimal traffic loss. Use the configuration command `set protocols evpn maintenance-mode erb-leaf action-type choice` to enable maintenance mode, and verify the status with `show evpn maintenance-mode status`. Ensure prechecks are validated to prevent disruptions, and manage the process efficiently with provided commands for deletion and validation.

[See [EVPN Maintenance Mode for Multihomed Leaf Isolation](#).]

- **EVPN multihoming and multitenancy support over colored IP fabric with BGP DPF (EX4100-24MP, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—You can leverage EVPN-VXLAN over colored IP fabric using BGP deterministic path forwarding (DPF) to support multihoming and multitenancy configurations for AI/ML applications. This functionality facilitates EVPN for Layer 3 networks with EVPN Type 5, enhancing network segmentation and resource allocation. By using a colored logical fabric, you can

achieve flexible routing as uncolored routes integrate seamlessly with all color-coded sessions, optimizing network efficiency and adaptability.

[See [BGP Deterministic Path Forwarding in a CLOS Network](#).]

- **Host route management in scaled, stretched VXLAN campus networks (EX4100-48MP, EX4100-24MP, EX4100-24T, EX4300-MP, 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)—**

We provide new routing policy options, sample routing policies, and new statements to optimize how remote host routes are managed across the access, distribution, and core layers in large-scaled VXLAN campus networks stretched across sites using DCI. With this solution, you can configure the network to install remote host routes in the core layer but not in the distribution and access layers. The distribution and access layer devices only install EVPN Type 5 subnet routes for remote host destinations. The configuration includes policies to ensure the EVPN Type 5 subnet routes are the preferred routes on the distribution and access layer devices. This design reduces the route table burden on access and distribution devices as the network scales.

[See [EVPN Host Route Management for Large-Scale Stretched Campus Networks](#).]

Interfaces

- **Support for transceiver hardening, mixed-speed uplinks, and 1000BASE-X mode on Falcon cores (EX2300-MP, EX2300-VC, EX4000-48MP, EX4100-48MP, EX4400-24MP, EX4400-24X, and EX4400-48F)—**Use hardened 10GbE SFP-T handling to reduce CPU usage by 30 percent and improve stability across restarts. Enhancements include:
 - Correct PHY initialization
 - TX disabled during initialization
 - Fewer auto-negotiation restarts
 - Optimized I2C locking
 - Link-state caching
 - Removed driver delays.

You can run 4x25 Gbps uplink modules in either uniform 25 Gbps or mixed 1 Gbps or 10 Gbps. Set the group speed for the chassis PIC 2 to 1 Gbps or 10 Gbps to enable mixed speed mode per transceiver. Set the speed to 25 Gbps for uniform 25 Gbps speed across every transceiver with that capability. The 4x10 Gbps modules operate in mixed mode always. The configuration is redundant, and the chassis PIC 2 operates in mixed-speed mode regardless of whether its speed is set to 1 Gbps or 10 Gbps, forming both 1 Gbps and 10 Gbps interfaces. 25 GbE and 10 GbE transceivers share the same SFP+ form factor, but 25 GbE transceivers are not supported in 10 Gbps ports. Inserting a 25

GbE transceiver into a 10 Gbps port does not create an interface. Always match transceiver speed to the port capability. For 1 Gbps fiber without SGMII, the device uses 1000BASE-X automatically.

[See [Port Speed on EX4400 Switches](#).]

Junos OS API and Scripting

- Support for libslax 3.1.6 and SLAX version 1.3 (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We've upgraded the libslax library to version 3.1.6, which corresponds to SLAX version 1.3. The upgraded library includes:
 - Slax processor enhancements including a new mode, additional options, and simplified argument parsing
 - New libslax extension library functions
 - Improved SLAX syntax options
 - New SLAX functions and enhancements to existing functions and statements
 - Hexadecimal numbers support in SLAX scripts

This update aligns SLAX processing with contemporary scripting conventions, ensuring efficient, readable scripting while maintaining backward compatibility.

[See [libslax Distribution Overview](#).]

Junos Telemetry

- Enhance VLAN/MAC statistics support for streaming telemetry (EX4300-MP, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48XP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, QFX5120-48Y, QFX5120-48Y-VC) You can enhance the granularity of streaming telemetry by subscribing to specific xpaths of leaf nodes using Yang state models generated from ODL files. This

subscription-based model provides precise and efficient data collection compared to the traditional RPC mechanism. Note that this solution supports only periodic streaming of telemetry data.

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

- **Sensor support for secure and dynamic packet capture (EX2300, EX2300-C, EX2300-MP, and EX2300-VC)** —EX2300, EX2300-C, EX2300-MP, and EX2300-VC devices support secure and dynamic packet capture. You can use this feature to capture packets from a device and send them over a secure channel to an external telemetry collector (in the cloud) for monitoring and analysis. Network professionals use real-time packet capture data to troubleshoot complex issues, including network and performance degradation, as well as poor end-user experience.

For secure packet capture, the maximum packet size that can be captured is 512 bytes, including the packet header and the data within. To use secure packet capture, include the `/junos/system/linecard/packet-capture` resource path using a Junos RPC call.

- For ingress packet capture, include the `packet-capture` option in the existing firewall filter configuration as follows:

```
[edit firewall family family-name filter filter-name term matchterm then packet-capture]
```

Use this configuration to send packet capture sensor data to the collector. Remove the packet-capture configuration after the data is sent to the collector. After the capture is done, ingress packets with the filter match conditions are trapped to the CPU. The trapped packets then go to the collector over a secure channel in JTI-specified format in key-value pairs through Remote Procedure Call (gRPC) transport.

- For egress packet capture on physical interfaces (ge-*, xe-*, mge-*, and et-*), include "packet-capture-telemetry," "egress," and "interface <interface-name>" at the [edit forwarding-options] hierarchy level.

For example:

```
set forwarding-options packet-capture-telemetry egress interface ge-0/0/0
```

```
set forwarding-options packet-capture-telemetry egress interface ge-0/0/10
```

You can add multiple interfaces on the device for egress packet capture. When configured, host-bound egress packets are captured from the interface and sent to the collector. As with the ingress configuration, remove the configuration when packet capture is not required.

For dynamic packet capture, subscribe to the resource path `/junos/system/linecard/packet-capture`. The device starts capturing the first "N" collector-bound packets for each physical interface present on the device when it transitions from the DOWN state to the UP state. The device then sends the packets securely to a collector.

By default, "N" is set as 50. For each interface, 50 ingress and 50 egress packets are captured. The data is captured from a Packet Forwarding Engine sensor and encoded using the Google Protocol Buffer format. The data is then sent over a secure channel using SSL encryption.

The collector receives the following packet attributes:

Table 3: Packet Attributes

Attribute Name	Description
total-length	The total length of the packet.
actual-length	The actual length of the packet.
packet-data	The packet data.
timestamp	Timestamp of the packet capture.
ifl-index	Logical interface index.
cos-queue	COS queue number.
direction	Indicates the direction of each packet.

For more information, see [Junos YANG Data Model Explorer](#) and [Supported gRPC and gNMI Sensors](#).

MAC Learning

- **Support to configure system timezone (EX Series)**—By default, mac-learning-logs use UTC timestamps. The logs appear in the output of the `show ethernet-switching mac-learning-log` command. Configure them to display in the system timezone for better alignment with local time. Use the `mac-learning-log system-timezone` option in the `l2-learning` command to display logs in the system timezone.

[See [l2-learning](#).]

Network Management and Monitoring

- **Ephemeral database default commit synchronize model changed to synchronous** (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-

VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)— We've changed the default commit synchronize model for the ephemeral database from the asynchronous model to the synchronous model. The synchronous model ensures better synchronization of ephemeral configurations across Routing Engines or Virtual Chassis members by processing commits synchronously. With this change, only the synchronous model supports synchronizing ephemeral data on devices that have graceful Routing Engine switchover (GRES) or nonstop active routing (NSR) enabled.

[See [Understanding Ephemeral Database Commit Synchronize Models](#).]

- **Adaptive sFlow polling intervals and buffer management improvements (EX Series switches)**— We've introduced dynamic sFlow polling interval and adaptive sampling configuration options to optimize performance and buffer management for sFlow enabled interfaces. You can adjust polling intervals based on interface counts, disable adaptive sampling, or set fixed polling intervals to maintain consistency despite interface changes. The default polling occurs every 12 seconds. If the interface count exceeds 100, intervals increase to 24 seconds. If the interface count is below 80, the intervals revert to 12 seconds. These enhancements improve efficiency in scaled environments, minimize scheduler slips, and ensure accurate network statistics.

To set a fixed polling interval and override dynamic adjustment, configure the `adaptive-interval interval value` option at `[set protocols sflow scale-mode]` hierarchy level.

To disable adaptive sampling and dynamic adjustment behaviors, use the configuration statements `set protocols sflow scale-mode adaptive-interval no-adaptive-sampling` and `set protocols sflow scale-mode disable`.

To view adaptive-sampling status and adaptive interval information, use the `show sflow` command.

[See [scale-mode](#), [sflow](#), and [show sflow](#).]

Post-Quantum Cryptography (PQC)

- **PQC signatures for software images with off-box verification (ACX Series, EX Series, MX Series, QFX Series, and SRX Series)**—Use post-quantum cryptography (PQC) signatures to ensure software images remain unaltered, protecting integrity and guarding against quantum threats. The images comply with algorithms recommended by Commercial National Security Algorithm 2.0 (CNSA 2.0):
 - ML-DSA-87 PQC algorithm for digital signatures
 - SHA-512 for hashing

For off-box verification, request the PQC signature from the support portal. Confirm the image's SHA-512 hash, retrieve the public key from the certificate, and validate the signature with your chosen verifier. PQC signatures provide additional security beyond existing legacy signatures.

[See [PQC Signatures for Software Images](#).]

Software Installation and Upgrade

- **Upgrade of firmware and software using a single command (EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, and EX4000-48T)**—You can streamline the upgrade process by integrating firmware upgrades with Junos OS software upgrades using a single command. During a software upgrade, the system stages the pending set for both the software and the firmware updates. After a reboot, the device operates with the latest versions.

To use this feature, specify the `update-firmware` option on either the `request system software add` command or the `request system software nonstop-upgrade` command.

[See [request system software add \(Junos OS\)](#) and [request system software nonstop-upgrade](#).]

Additional Features

We've extended support for the following features to these platforms.

- **Supported transceivers, optical interfaces, and DAC cables.** (EX4100, EX4400, EX4400-24X, EX4400-EM-4Y, EX4650, MX304, MX10004, MX10008, QFX5120, and SRX4700). Select your product in the [Hardware Compatibility Tool](#) 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 optical transceiver becomes available.
- **Renaming OpenSSH implementation to JSSH (all platforms)**—The OpenSSH implementation in Junos OS is renamed "JSSH" to highlight its customized nature.

[See [ssh](#).]

What's Changed

IN THIS SECTION

- [Dynamic Host Configuration Protocol | 30](#)
- [General Routing | 30](#)
- [Interfaces | 32](#)
- [MPLS | 32](#)
- [Network Management and Monitoring | 32](#)
- [Platform and Infrastructure | 33](#)

- [Routing Protocols | 33](#)
- [Subscriber Access Management | 34](#)
- [User Interface and Configuration | 34](#)

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



NOTE: The EX2300 and EX3400 models are documented but not supported in this release.

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

General Routing

- A new counter "Sessions hit due to high rate" is added to `show services service-sets screen-session-limit-counters` command for all subscriber traffic. This counter tracks the sessions that come up on the screen irrespective of the alarm-without-drop configuration. When alarm-without-drop option is disabled, all the counters display updated statistics. When alarm-without-drop is enabled, then, the screen-drop counters on `show services service-sets statistic screen-drop` command do not increase. The "sessions hit due to high rate" value is displayed.

[See [alarm-without-drop \(IDS Screen Next Gen Services\)](#), [show services service-sets statistic screen-drops \(Next Gen Services\)](#), and [show services service-sets statistic screen-session-limit-counters \(Next Gen Services\)](#).]PR1849594

- When you run the `request vmhost zeroize` command to zeroize a single Routing Engine on a dual Routing Engine device, the CLI incorrectly displays a message indicating that it will zeroize both Routing Engines. [PR1869854](#)
- Validation of `filepath/vm-primary(filepath)` mount during JDM installation or upgrade (Junos Node Slicing External Server Deployment)**—During installation or upgrade of the Juniper Device Manager (JDM) on external servers running Junos Node Slicing Release 25.2 or later, an issue occurred with the validation of the `filepath/vm-primaryfilepath` mount that stores the GNF images. When

filepath/vm-primary filepath was mounted using logical volumes (LVM), JDM would fail to detect that the underlying storage was an SSD. This issue is now fixed. However, the fix introduces a new dependency on the LVM2 package in the host OS. This package is included by default in standard installations of both RHEL and Ubuntu external servers. However, it is advised that you check if the LVM2 package is already installed on the host before installing or upgrading JDM.[PR1877593](#)

- On the MPC7E-10G line card, when you configure the 10-Gigabit Ethernet ports to operate as 1-Gigabit Ethernet ports, use the speed statement at both the [edit interfaces *interface name* together-options] and edit interfaces *interface name* hierarchy levels. [PR1879198](#)
- **Control Maximum 802.1X Client Connections per Interface**—By default, dot1x interfaces configured in multiple supplicant mode have a client limit of 100 authenticated connections per interface. Any additional connection attempts beyond this limit will be automatically blocked.
- **New option for debug collector data storage path** We've included the option `outdir` to specify an output directory for storing debug collector data in a customised path. This allows you to organise and access diagnostic information more efficiently, adapting storage to your specific requirements. [PR1889710](#)
- **Correct auto-negotiation status display for 1G full-duplex links (EX2300-48MP, EX3400-48T, EX4100-48MP, EX4100-48P, and EX4400-48P)**—You should be aware that when operating at a speed of 1G with full-duplex mode and no auto-negotiation, the system internally advertises auto-negotiation capabilities despite displaying incorrect status as "disabled" and "No-auto-negotiation" in interface commands. You need to remove the 1G option for copper ports while configuring speed to avoid this discrepancy.
- Topic updated with limitations on hop-limit match condition for EX4100 and EX4400. [PR1890156](#)
- **Static flowspec prefix validation**—Commit validation now rejects static flowspec destination and source prefixes that use a shorter mask than the configured address family permits. The configuration must use a valid prefix length before the commit can succeed. This alignment with static route validation prevents invalid flowspec entries from being installed and ensures consistent routing policy enforcement across IPv4 and IPv6 configurations.[PR1903233](#)
- **Enhanced ping command XML output (Junos OS and Junos OS Evolved)**—The XML output for the CE-IP ping has enhanced details to differentiate between successful and failed ping attempts. The update introduces new tags that track the total number of packets sent and received, while also highlighting host unreachable replies and the overall packet loss percentage [See [Supported and Unsupported Features for CE-IP Ping](#).][PR1906967](#)
- Support added for new payload match conditions (EX4100 , EX4400 , EX5200, and QFX5120)- Support added for payload-source-ipv4-address, payload-destination-ipv4-address, payload-source-ipv6-address, payload-destination-ipv6-address, payload-source-mac-address, and payload-destination-mac-address match conditions for IPv4 and IPv6 traffic which can be applied on the ingress port only.

[See [Firewall Filter Match Conditions for IPv4 Traffic](#).]PR1914083

- Topic updated with Platform Specific Behavior section - [Firewall Filter Flexible Match Conditions](#). PR1917415
- `request jnu role delete` **Command Unhidden (Junos OS and Junos OS Evolved)**—The `request jnu role delete` command, which was previously hidden, is now available for use. This command allows users to delete the role of controllers or satellites directly through the CLI. This enhancement improves administrative control and flexibility in role management within the network infrastructure.

[PR1930291](#)

Interfaces

- **Mixed 1G and 10G speed support on 4x25G ULM module (EX4400)**—Support has been added for operating 1G and 10G interface speeds concurrently on the 4x25G ULM module when it is configured in 1G/10G mode. When you configure the quad using `set chassis fpc 0 pic 2 port 0 speed 10g` or `set chassis fpc 0 pic 2 port 0 speed 1g`, the system detects and brings up both 1G and 10G transceivers in that quad, creates interfaces for 1G and 10G transceivers, and skips creating interfaces for any 25G transceivers present in the quad. This enhancement simplifies interface management and increases flexibility by allowing you to mix 1G and 10G optics within the same 4x25G ULM quad under a single speed configuration.

[See [Port Speed on EX4400 Switches](#).]PR1969800

MPLS

- Label field naming update in `show route forwarding-table extensive` output (Junos OS and Junos OS Evolved)—The CLI output for MPLS forwarding entries now uses "Label" instead of "Load balance label." This name change standardizes the field name and improves overall clarity when reviewing MPLS forwarding information. The change affects only the output text and does not modify label allocation, forwarding behavior, or traffic handling. No configuration changes are required.

[PR1899543](#)

Network Management and Monitoring

- Ephemeral database default commit synchronize model changed to synchronous (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-48MP, EX4000-24P, EX4000-24T, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-

VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, and MX10008, MX10016)—We've changed the default commit synchronize model for the ephemeral database from the asynchronous model to the synchronous model. With this change, we've deprecated the `allow-commit-synchronize-with-gres` statement and only the synchronous model supports synchronizing ephemeral data on devices that have graceful Routing Engine switchover (GRES) or nonstop active routing (NSR) enabled.

[See [Understanding Ephemeral Database Commit Synchronize Models](#).]

- **Shell Command Logging Enhancement**--All shell commands executed on the device and root sessions are now logged by default. This enhancement ensures enhanced security and auditability by capturing all commands entered in any shell environment, preventing bypassing audit logging.[PR1867216](#)
- **Deprecation of shell option**--The `shell` option no longer requires a separate configuration and is now the default behavior. Deprecating the shell option enhances efficiency and simplifies management tasks.

Platform and Infrastructure

- **Commit validation for unique user IDs**--We have added support to validate the user configuration to ensure that each user is assigned a unique UID. A commit fails if duplicate UIDs are detected, ensuring stronger validation and preventing identity conflicts. Previously, a commit was successful even when multiple users shared the same UID, triggering only a warning and logging a syslog message.[PR1890656](#)

Routing Protocols

- **SRv6 MSD advertisement compliance**--MSD type 43 (SRH Max T.insert) is no longer advertised in IS-IS SRv6 node capability TLVs to match the latest standards draft. This alignment maintains interoperability with peer devices and tools that expect the reduced MSD type list and ensures SRH limits are reported consistently across the network.[PR1879051](#)
- **Auto BGP peer discovery for unnumbered interfaces of iBGP peers**--Junos supports single-hop iBGP sessions for unnumbered peering. Previously, the iBGP implementation created an indirect next-hop for route resolution, even if the iBGP peer was directly connected.[PR1884861](#)
- **BGP route-refresh warning messages added**--New messages have been introduced to the `show bgp diagnostics warnings` command output. These messages provide diagnostic information when the number of BGP route-refresh messages exceed predetermined thresholds within specific time intervals of 5 minutes, 30 minutes, and 1 hour.

[See [show bgp diagnostics warnings](#).][PR1909467](#)

- **BGP UPDATE messages must have the marker field set to all ones**—The marker field in BGP header is a 16-octet field included for compatibility that must be set to all ones. All Junos OS and Junos OS Evolved devices send a notification and tear down the BGP session if the marker field in the header of the received update message is not set to all ones.[PR1955678](#)

Subscriber Access Management

- **Addition of message-authenticator and no-message-authenticator attributes under access radius-server, access profile radius-server, and system radius-server hierarchies**—set message-authenticator if you require the RADIUS server to include the message authenticator attribute in replies to Access-Request messages. set no-message-authenticator to not require that attribute.[PR1871147](#)

User Interface and Configuration

- **Updated Annotate Command Behavior**—The annotate command now correctly handles multi-line comments that start with `#`. This change ensures that all lines of a multi-line comment remain comments, preventing unauthorized configuration changes when the configuration is committed, reloaded, or rolled back. These improvements enhance security by mitigating potential privilege escalation attacks from users with limited configuration editing permissions.[PR1868636](#)
- **Enhanced Permission Checks for Rename/Copy Operations**—New permission checks have been introduced for rename and copy operations within the configuration hierarchy. These checks parse the hierarchy being modified to ensure the user has the required permissions for the hierarchy and its sub-hierarchies. If permissions are insufficient, the operation will fail and a "Permission denied" error will be displayed. This enhancement ensures that configuration modifications are performed only by authorized users, improving security and preventing unintended changes to critical system settings.[PR1882303](#)
- **Enhanced file archive exclusion capability (PTX10008)**—The file archive operational command now supports exclude option where you can specify a folder or regular expression to exclude from final archive. (Example, exclude traces from `/var/log`). This enhancement improves the flexibility and efficiency of the archiving process, allowing you to manage log files more effectively. To exclude folders, you can use one of the below formats:

```
user@router> file archive source /var/log destination /var/tmp/abc exclude jam* exclude wtmp*
exclude ksyncd exclude aaasd
user@router> file archive exclude [license* shmlog ksyncd] source /var/log/
destination /var/tmp/abc
```

[PR1884901](#)

- **Stale ui-state.db data in persistent NETCONF sessions post-mgd restart**—Existing NETCONF sessions might fetch stale data from ui-state.db after mgd -N restart. New sessions correctly map the

refreshed database. Scripts must establish new sessions post-restart to access updated values. Functional configuration remains unaffected. [Script failures monitoring "local-host" NETCONF sessions]-Scripts might fail when including "local-host" NETCONF sessions in monitoring operations. Internal sessions are now excluded from tracking. Scripts must filter out "local-host" sessions. No impact to internal application functionality.[PR1888557](#)

- **Generate genstate YANG modules on Junos devices**—You can use the `show system schema operational` command or equivalent RPC to generate the genstate YANG modules in the specified output directory on a device.

[See [show system schema](#).][PR1910388](#)

Known Limitations

IN THIS SECTION

- [General Routing | 35](#)
- [Infrastructure | 36](#)

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



NOTE: The EX2300 and EX3400 models are documented but not supported in this release.

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 EX4400-48F, whenever JOJI of optics is performed, the below messages are expected BCM Error: `API bcm_plp_seahawks_module_read(phy_info, addr, offset, size, data)` at `tvb_bcm_seahawks_eeprom_read:952 -> -8`[PR1742075](#)
- Ex-Hardening:Local/Remote fault insertion from TG is failing. [PR1789999](#)
- The issue is specific to the arm based platforms(EX4100, EX4000). This log is triggered under certain scenarios of DHCP failure events such as `AS_PKT_DAI_FAILED`, `AS_PKT_DHCP_DROPPED`, etc. As a

workaround, we enabled one more log with the proper values under system events. For now, user can consider the log with proper values. [PR1890822](#)

- On MX AFT based platforms, packets injected over IRB may get dropped in the PFE with an 'MTU exceeded' exception if the underlying Layer 2 interface is in trunk mode. [PR1901133](#)

Infrastructure

- When upgrading from releases before Junos OS Release 21.2 to Release 21.2 and onward, validation and upgrade might fail. The upgrade requires using the 'no-validate' option to complete successfully. <https://kb.juniper.net/TSB18251>[PR1568757](#)

Open Issues

IN THIS SECTION

- [General Routing | 36](#)
- [Platform and Infrastructure | 37](#)
- [Routing Protocols | 37](#)

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



NOTE: The EX2300 and EX3400 models are documented but not supported in this release.

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

General Routing

- EX4100-24mp,48mp,24p/t,48p/t,F-24p/t,F-48-p/t: In an interop scenario, when using 1G SFP Optic on PIC-2, auto-negotiation should be disabled on the peer. [PR1657766](#)
- EX4300MP: VC member status toggling between "Inactive" and "NotPrsnt" state after member downgrade. [PR1751871](#)

- Time Domain Reflectometry (TDR) support for detecting cable breaks and shorts aborts intermittently on some random ports.[PR1820086](#)
- On EX4400 devices with 4x25G or 1x100g ULM, when we perform PIC online, CPU overuse by CMQFX thread might be seen for as much as 3.5seconds. [PR1870962](#)
- With latest changes in 25.2R1 on EX4000 and EX4100, after bootup there are error logs seen in syslog as below: fpc0 brcm_bcm_l3_ingress_create:52Could not create l3 ingress for error Operation disabled. dc-pfe[11749]: brcm_bcm_l3_ingress_create:52Could not create l3 ingress for error Invalid parameter. These errors have no functionality impact. Request to ignore these logs seen just after bootup. [PR1872146](#)
- When a non-PoE SKU is configured as the master and PoE SKUs are configured as members, the PoE status will be "disabled" in cli even though PDs are powered up. However, if the device is rebooted while in this state, PoE functionality will not resume afterward. [PR1881788](#)
- In the EX4400-48F systems, a 10G-BaseT transceiver that was earlier up may not come up post a reboot/image upgrade event; The transceiver might go undetected causing the interface to not be created in the system.[PR1887303](#)
- EX4400-48F: When we have SFP-SX optics plugged in EX4400-48F device in the ports 0 to 35 and it is rebooted, the activity LED remains on. [PR1899248](#)
- With SFP-T optics and DUT configured speed=100M with Auto-Neg enabled and peer configured speed=10M with Auto-Neg enabled, link might not come up. [PR1905956](#)
- request system firmware upgrade jfirmware allfw only upgrading CPLD not upgrading Uboot[PR1942521](#)

Platform and Infrastructure

- In a rare scenario, due to timing issues, the Packet Forwarding Engine (PFE) crash is observed on Junos EX4300 platforms. This causes traffic loss until the PFE comes up. [PR1720219](#)
- It is noticed that EX4300 switches after an upgrade of Junos from 21.2R3-SX to 21.4R3-SX may exhibit a higher Cpu. Issue is resulting from fast path thread profiling code. It takes on an average 1 ms more for one fast path thread cycle, cumulatively overall fast path thread usage had increased. Thread profiling code has been optimised and the issue is fixed in the future Junos.[PR1794342](#)

Routing Protocols

- Traffic convergence issue was resolved with the work around provided. BFD issue was known limitation. [PR1888400](#)

Resolved Issues

IN THIS SECTION

- [General Routing | 38](#)
- [EVPN | 40](#)
- [Forwarding and Sampling | 40](#)
- [J-Web | 40](#)
- [Network Management and Monitoring | 40](#)
- [Routing Protocols | 40](#)
- [User Interface and Configuration | 40](#)

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



NOTE: The EX2300 and EX3400 models are documented but not supported in this release.

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

General Routing

- IS-IS Level 2 disabled after upgrade due to OpenConfig YANG model change. [PR1798178](#)
- The dcpfe process crashes on specific Junos QFX and EX platforms due to memory corruption. [PR1856424](#)
- Lower rate is not negotiated even if auto-negotiation is configured on EX3400 platforms. [PR1858406](#)
- EX4400: Error message 'Failed to read eeprom' intermittently on SFP-T. [PR1858986](#)
- The 'fxpc' process utilization shows above 80% with 100G AOC/DAC or 100G optics connected on interface with FEC mode mismatch. [PR1860519](#)
- On Junos QFX5000 series and EX4000 series platforms, an fxpc process crash triggers an FPC reboot. [PR1866815](#)

- The PTP packets are dropped when IGMP snooping is enabled [PR1873129](#)
- On EX4100/EX4400s platforms PoE powered devices connected do not come up when adding a second power supply unit. [PR1876675](#)
- There is a PoE short circuit alarm after upgrading the device. [PR1879702](#)
- JMA package fails to initialise after a power cycle on EX4650/QFX-5E series devices. [PR1882472](#)
- Link-mode shows Half-Duplex when interface is set to 100M/10M with full-duplex and no-auto-negotiation. [PR1883130](#)
- Error messages are seen when the device is loaded with baseline EVPN-VxLAN configuration [PR1884063](#)
- snmp mib walk DomCurrentLaneAlarms does not show proper lanes Values in the latest builds [PR1886889](#)
- Configuring loopback firewall optimization causes IPv6 BGP session down. [PR1886981](#)
- JTASK_IO_CONNECT_FAILED messages are seen on the device. [PR1887336](#)
- On some EX or QFX platforms, 25G interfaces remain down with a default FEC(Forward Error Correction) 74 value when peer device has default FEC as None. [PR1890164](#)
- The pkid crash is observed during enrolment of device's local certificate through SCEP. [PR1892297](#)
- FXPC crash on EX4k VC during GRES operation. [PR1893742](#)
- Physical Interface device remains down by continuous system reboots on EX and QFX5K platforms. [PR1894136](#)
- The enable-pxe-boot is not working as expected on specific EX or QFX platforms with Easy EVPN LAG single home server topology. [PR1895433](#)
- EX2300/EX3400: TFTP installation failure. [PR1900881](#)
- [ex4000]:: Dot1x client remains in server-fail vlan post radius-reachability timer expiry. [PR1901304](#)
- PFE crashes on QFX5K and EX4K platforms in standalone and VC setups due to logging of PFE. [PR1901837](#)
- Intermittent kernel panic results in device reboot or fxpc crash. [PR1902609](#)
- Significant system slowness is observed post software upgrade on QFX5120-48YM. [PR1902869](#)
- During virtual chassis switchover causes default dead route creation. [PR1904091](#)
- Stale entry in MAC-IP table affects ARP resolution. [PR1905196](#)

- EX3400 EX4400-48F: Auto-negotiation is not displayed in 'show interfaces' command output. [PR1909608](#)
- On-change telemetry events dropped when l2ald telemetry queue memory limit is not configured. [PR1914423](#)
- snmp mib walk DomCurrentLaneAlarms does not show proper lanes Values in the latest builds. [PR1886889](#)

EVPN

- EVPN routes take longer to install into the FIB after being learned via BGP. [PR1893671](#)

Forwarding and Sampling

- Filter name mapping error in routing instances with attached firewall filter. [PR1889028](#)
- Memory leak in FPC during login/logout. [PR1890097](#)

J-Web

- Junos OS: J-Web: Multiple vulnerabilities resolved in PHP software (CVE-2023-0567, CVE-2023-0662, CVE-2023-3823, CVE-2023-3824, CVE-2023-0568). [PR1725808](#)

Network Management and Monitoring

- Syslog forwarding intermittently stops post DUT reboot on virtual devices. [PR1853209](#)

Routing Protocols

- BFD sessions will not come up on Junos OS and Junos OS Evolved platforms due to keychain names overlapping. [PR1912250](#)

User Interface and Configuration

- The mgd process crash is seen on all Junos and Junos Evolved platforms when FQDN is configured along with ephemeral database. [PR1878430](#)
- J-Web EX app-package upload fails via web interface on EX series running Junos 25.1/25.2. [PR1887595](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 41

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 4: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
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 | 42](#)
- [What's Changed | 43](#)
- [Known Limitations | 43](#)
- [Open Issues | 43](#)
- [Resolved Issues | 44](#)
- [Migration, Upgrade, and Downgrade Instructions | 44](#)

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

- [Dynamic Host Configuration Protocol](#) | 43

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

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

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 open 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](#) | 44

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 5: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
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 Juniper Secure Connect

IN THIS SECTION

- [What's New | 46](#)
- [What's Changed | 46](#)
- [Known Limitations | 46](#)
- [Open Issues | 46](#)
- [Resolved Issues | 47](#)

What's New

There are no new features or enhancements to existing features in this release for Juniper Secure Connect.

What's Changed

IN THIS SECTION

- [Dynamic Host Configuration Protocol](#) | 46

Learn about what changed in this release for Juniper Secure Connect.

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

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 MX Series

IN THIS SECTION

- [What's New in 25.4R1-S1 | 47](#)
- [What's New in 25.4R1 | 58](#)
- [What's Changed | 84](#)
- [Known Limitations | 91](#)
- [Open Issues | 92](#)
- [Resolved Issues | 99](#)
- [Migration, Upgrade, and Downgrade Instructions | 108](#)

What's New in 25.4R1-S1

IN THIS SECTION

- [Hardware | 48](#)
- [MACsec | 57](#)

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

Hardware

- **New MX301 Router (MX Series)**—The Juniper Networks® MX301 Universal Routing Platform is a next-generation fixed configuration router in a compact and power efficient 1-U form factor. The router provides flexible connectivity options with 26 WAN ports that offer port configurations from 1GbE up through 400GbE, delivering speeds of up to 1.6 Tbps. The MX301 supports six fan modules with front-to-back airflow and two 850-W AC, DC, or HVAC/DC PSUs.

The MX301 is ideal for high-density, high-throughput networking deployments in enterprise, cloud, service provider edge, and AI edge clustering use cases.

To install the MX301 Router and perform initial configuration, routine maintenance, and troubleshooting, see the [MX301 Universal Routing Platform Hardware Guide](#). For a detailed list of available features, see [Feature Explorer](#).

Table 6: MX301 Feature Support

Feature	Description
Class of Service	• Support for Class of Service (CoS) features including: • Behaviour Aggregate classifiers and rewrites • Hierarchical Class of Service (HCoS)/Per VLAN with a queuing delay bandwidth buffer of 100ms for interfaces with 1G or higher speeds, and a buffer of 500ms for interfaces with less than 1G speeds. • Trio QoS model with five guaranteed regions and four excess regions for the scheduler (in both normal and enhanced priority modes) • Eight queues with strict prioritization [See Behavior Aggregate Classifiers, Hierarchical Class of Service Overview, and Class of Service for Routers.]

Table 6: MX301 Feature Support *(Continued)*

Feature	Description
Dynamic Host Configuration Protocol	• Support for Dynamic Host Configuration Protocol (DHCP) v4/v6 in non-broadband Edge solutions, including: • Local server • Dynamic reconfiguration of clients from a DHCP local server • Smart, stateful, and forward-only (stateless) relay • Snooping • Integrated routing and bridging (IRB) • EVPN-VXLAN support for centrally routed bridging (CRB) access models only. We don't support EVPN-VXLAN on edge-routed bridging (ERB) access models. [See DHCP Relay Agent, DHCP Subscriber Access Networks Overview, DHCP Snooping for Network Security, DHCP Smart Relay in EVPN-VXLAN, and Dynamic Client/Server Interaction for DHCPv6.]
High availability and resiliency	• Support for platform resiliency, line card, and chassis resiliency. Routing Engine resiliency runs on Junos OS software and covers the basic chassis management including Routing Engine, fans, and power supply modules (PSMs). ULC resiliency runs on the Junos OS Evolved software including the YT ASIC, Packet Forwarding Engine, PFE FPGA, and the Ethernet switch. [See Platform Resiliency.] • Support for Routing and Control Board (RCB) and Packet Forwarding Engine resiliency. [See show chassis errors active.]

Table 6: MX301 Feature Support (*Continued*)

Feature	Description
Interfaces	- Support for optics temperature sensors in the Junos Environment Monitoring (EM) policy, including: - Polling temperature readings of optical modules in the system periodically to automatically manage the fan speed - Triggering optics shutdown for 50GbE, 100GbE, and 400GbE optics firmware when the high-temperature threshold is breached The EM policy is enabled by default on all 50GbE, 100GbE, and 400GbE optics interfaces, except for loopback optics and direct attach copper (DAC) cables. You can use the <code>set chassis fpc <i>fpc_slot</i> pic <i>pic_slot</i> port <i>port_no</i> 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.] - Select your product in the https://apps.juniper.net/hct/product (HCT) 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.
Junos telemetry interface (JTI)	Support for JTI platform sensors. A basic configuration is sufficient to enable a gRPC/gNMI subscription. [See gRPC Services for Junos Telemetry Interface and Guidelines for Telemetry Data Subscriptions over gNMI.]

Table 6: MX301 Feature Support (*Continued*)

Feature	Description
Layer 2 features	• Support for Layer 2 (L2) features, including: • Layer 2 protocol tunneling (L2PT) • Virtual private wire service (VPWS) • EVPN-MPLS and EVPN-VXLAN overlay networks • EVPN routing policies • Multiple VLAN Registration Protocol (MVRP) and Ethernet ring protection switching (ERPS) • VPLS over transport class tunnels • Headend termination of pseudowire services in a VPLS-enabled virtual switch • Q-in-Q tunneling • Basic DHCP snooping [See Layer 2 Protocol Tunneling, Overview of Headend Termination for EVPN VPWS, Multicast Support in EVPN-VXLAN Overlay Networks, Routing policies for EVPN, Understanding Multiple VLAN registration protocol (MVRP) for Dynamic VLAN Registration, Ethernet Ring Protection Switching Overview, Introduction to VPLS, Layer 2 Services on Pseudowire Service Interface Overview, Configuring Q-in-Q Tunneling, and DHCP Snooping.]

Table 6: MX301 Feature Support (*Continued*)

Feature	Description
Layer 3 features	• Support for Layer 3 (L3) features, including: • Network slicing with routing-based and firewall-based slice steering • Symmetric consistent hashing for symmetric load-balancing with correlated source IP and destination IP load-balancing hash-key (both forward and reverse direction) • OpenConfig-compliant Segment Routing-Traffic Engineering (SR-TE) sensor compliance • Adaptive load balancing (ALB) on aggregate Ethernet (AE) interfaces • Consistent hashing on flex-algo routes in a BGP network • Transit statistics for pseudowire and PWHT interfaces • GRE tunnels over PWHT interfaces • Seamless MPLS L3 support with pseudowire subscriber over redundant logical tunnel and pseudowire subscriber over logical tunnel. • Multicast routing protocols • MPLS Segment routing • Layer 3 (L3) unicast forwarding and GRE tunneling • RFC-2544-based benchmarking tests [See Hierarchical Class of Service for Network Slicing, Actions in Routing Policy Terms, YANG Data Model Explorer, Understanding Aggregated Ethernet Load Balancing, show interfaces extensive, MPLS Pseudowire Subscriber Logical Interfaces, MPLS Overview, Multicast Routing Protocols, Generic Routing Encapsulation (GRE), and Understanding RFC 2544-Based Benchmarking Tests.]

Table 6: MX301 Feature Support (*Continued*)

Feature	Description
MACsec	Support for MACsec bounded delay protection. Enable Media Access Control Security (MACsec) bounded delay protection to protect your network against man-in-the-middle attacks. When enabled, the device ensures that a frame is not sent after a delay of two seconds or more. MACsec periodically compares the number of frames transmitted to the number of frames received. If a frame is sent but not received within two seconds, such as during a man-in-the-middle-attack, the packet is dropped. [See bounded-delay.]
Network management and monitoring	Support for port mirroring. Analyzer is currently not supported. [See Port Mirroring and Analyzers.]
Protection against DDoS attacks	Support for distributed denial-of-service (DDoS) protection. DDoS protection is enabled by default. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]
Routing protocols	Support for L3 routing protocols. [See Routing Protocol Process Overview.]
Routing policy and firewall filters	Support for forwarding firewalls and 1GbE firewall support. Fabric loopback will be replaced by WAN stream loopback, which can impact mainline throughput. [See Understanding Firewall Filter Match Conditions and Service Filter Overview.]

Table 6: MX301 Feature Support *(Continued)*

Feature	Description
Security services	• Junos OS supports IKE to define mechanisms for key generation and exchange, and manage security associations (SAs). Inline IPsec is supported in Unified Services Framework (USF) mode. The IPsec architecture provides a security suite for the IPv4 and IPv6 network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. ESP in tunnel mode is supported. [See IPsec Overview and Enabling and Disabling Next Gen Services.]

Table 6: MX301 Feature Support (*Continued*)

Feature	Description
Services applications	• Support for inline services, including: • Inline NAT—NAT44 and NPTv6 • Inline softwires—Mapping of Address and Port with Encapsulation (MAP-E) • Mapping of Address and Port with Translation (MAP-T) • IPv6 rapid deployment (6rd) [See Inline NAT, Configuring Mapping of Address and Port with Encapsulation (MAP-E), and Configuring Inline 6rd.] • FlowTapLite support. [See Understanding the FlowTap and FlowTapLite Services.] • Inline video monitoring support, including: • Inline video monitoring using media delivery indexing (MDI) • Inline video monitoring of IPv4-over-MPLS flows to measure media delivery index (MDI) metrics • Inline video monitoring of IPv6 flows and IPv6-over-MPLS flows to measure media delivery index (MDI) metrics [See Understanding Inline Video Monitoring on MX Series Routers, video-monitoring, and show services video-monitoring mdi flows fpc-slot.] • Inline monitoring services support to provide the flexibility to monitor different streams of traffic at different sampling rates on the same interface. [See Inline Monitoring Services Configuration.] • Inline active flow monitoring support, including: • EVPN traffic

Table 6: MX301 Feature Support (*Continued*)

Feature	Description
	• MPLS, MPLS-IPv4, and MPLS-IPv6 traffic • IPv4 or IPv6 traffic, including traffic on next-hop based GRE tunnels and ps interfaces Both IPFIX and V9 formats are supported. [See Understand Inline Active Flow Monitoring.]
Subscriber management and services	• Support for subscriber services in Broadband Edge (BBE) and Broadband Network Gateway (BNG) solutions. Inline carrier-grade NAT (CGNAT) is currently not supported. [See Broadband Subscriber Services.] • Support for broadband network gateway (BNG) control and user plane separation (CUPS) architecture. Juniper BNG CUPS is a cloud application that separates the BNG control plane components from the native Junos OS. [See Juniper BNG CUPS Overview.] • Support for 5G Access Gateway Function (AGF). [See Access Gateway Function.] • Support for 5G Wireless CUPS user plane function (UPF). The UPF, supported on MX Series routers, handles all user data traffic according to the CUPS model. Junos Multi-Access UPF supports MX Series routers functioning as UPFs and provides high-throughput 5G fixed wireless and mobile wireless service in non-standalone (NSA) mode. [See MX Series Router As Junos Multi-Access User Plane.]

Table 6: MX301 Feature Support (*Continued*)

Feature	Description
Software installation and upgrade	- Secure boot and common BIOS support. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms. [See Junos OS Overview.] - Firmware upgrade support. [See request system firmware upgrade.] - Support for secure zero-touch provisioning (SZTP). Use RFC 8572-based SZTP to bootstrap your remotely located network devices that are in a factory-default state. SZTP enables mutual authentication between the bootstrap server and the network device before the remote network device is accessed for initiating zero-touch provisioning. [See Secure Zero Touch Provisioning.]

MACsec

- Support for a custom EAPoL EtherType to improve network tunneling of MACsec packets (MX301)**—MACsec uses Extensible Authentication Protocol over LAN (EAPoL) as a transport protocol to establish sessions. Some networks filter packets based on their EtherType value. By default, the EtherType for all EAPoL packets is 0x888e. To ensure the network tunnels the MACsec packets properly, you can set a custom EtherType for EAPoL packets.

To configure an EAPoL profile with a custom EtherType, use the `ether-type ether-type-value` statement at the `[edit forwarding-options custom-eapol-ether-type-profiles eapol-profile-name]` hierarchy level. To apply the EtherType to MACsec packets, configure the `eapol-ethertype-profile eapol-profile-name` statement at the `[edit security macsec connectivity-association ca-name mka]` hierarchy level.

[See [Media Access Control Security \(MACsec\) over WAN](#), [custom-eapol-ethertype-profiles](#), and [mka](#).]

What's New in 25.4R1

IN THIS SECTION

- [Hardware | 58](#)
- [Chassis | 68](#)
- [EVPN | 68](#)
- [High Availability | 69](#)
- [Interfaces | 70](#)
- [Junos OS API and Scripting | 70](#)
- [Precision Time Protocol \(PTP\) | 71](#)
- [Junos Telemetry | 73](#)
- [MACsec | 75](#)
- [MPLS | 75](#)
- [Network Management and Monitoring | 76](#)
- [NextGen Port Extender \(NGPE\) | 77](#)
- [Post-Quantum Cryptography \(PQC\) | 77](#)
- [Routing Protocols | 77](#)
- [Services Applications | 79](#)
- [Source Packet Routing in Networking \(SPRING\) or Segment Routing | 80](#)
- [Subscriber Management and Services | 82](#)
- [Additional Features | 83](#)

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

Hardware

- **New MX301 Router (MX Series)**—The Juniper Networks® MX301 Universal Routing Platform is a next-generation fixed configuration router in a compact and power efficient 1-U form factor. The router provides flexible connectivity options with 26 WAN ports that offer port configurations from 1GbE up through 400GbE, delivering speeds of up to 1.6 Tbps. The MX301 supports six fan modules with front-to-back airflow and two 850-W AC, DC, or HVAC/DC PSUs.

The MX301 is ideal for high-density, high-throughput networking deployments in enterprise, cloud, service provider edge, and AI edge clustering use cases.

To install the MX301 Router and perform initial configuration, routine maintenance, and troubleshooting, see the [MX301 Universal Routing Platform Hardware Guide](#). For a detailed list of available features, see [Feature Explorer](#).

Table 7: MX301 Feature Support

Feature	Description
Class of Service	• Support for Class of Service (CoS) features including: • Behaviour Aggregate classifiers and rewrites • Hierarchical Class of Service (HCoS)/Per VLAN with a queuing delay bandwidth buffer of 100ms for interfaces with 1G or higher speeds, and a buffer of 500ms for interfaces with less than 1G speeds. • Trio QoS model with five guaranteed regions and four excess regions for the scheduler (in both normal and enhanced priority modes) • Eight queues with strict prioritization [See Behavior Aggregate Classifiers, Hierarchical Class of Service Overview, and Class of Service for Routers.]

Table 7: MX301 Feature Support (*Continued*)

Feature	Description
Dynamic Host Configuration Protocol	• Support for Dynamic Host Configuration Protocol (DHCP) v4/v6 in non-broadband Edge solutions, including: • Local server • Dynamic reconfiguration of clients from a DHCP local server • Smart, stateful, and forward-only (stateless) relay • Snooping • Integrated routing and bridging (IRB) • EVPN-VXLAN support for centrally routed bridging (CRB) access models only. We don't support EVPN-VXLAN on edge-routed bridging (ERB) access models. [See DHCP Relay Agent, DHCP Subscriber Access Networks Overview, DHCP Snooping for Network Security, DHCP Smart Relay in EVPN-VXLAN, and Dynamic Client/Server Interaction for DHCPv6.]
High availability and resiliency	• Support for platform resiliency, line card, and chassis resiliency. Routing Engine resiliency runs on Junos OS software and covers the basic chassis management including Routing Engine, fans, and power supply modules (PSMs). ULC resiliency runs on the Junos OS Evolved software including the YT ASIC, Packet Forwarding Engine, PFE FPGA, and the Ethernet switch. [See Platform Resiliency.] • Support for Routing and Control Board (RCB) and Packet Forwarding Engine resiliency. [See show chassis errors active.]

Table 7: MX301 Feature Support (*Continued*)

Feature	Description
Interfaces	- Support for optics temperature sensors in the Junos Environment Monitoring (EM) policy, including: - Polling temperature readings of optical modules in the system periodically to automatically manage the fan speed - Triggering optics shutdown for 50GbE, 100GbE, and 400GbE optics firmware when the high-temperature threshold is breached The EM policy is enabled by default on all 50GbE, 100GbE, and 400GbE optics interfaces, except for loopback optics and direct attach copper (DAC) cables. 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.] - Select your product in the https://apps.juniper.net/hct/product (HCT) 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.
Junos telemetry interface (JTI)	Support for JTI platform sensors. A basic configuration is sufficient to enable a gRPC/gNMI subscription. [See gRPC Services for Junos Telemetry Interface and Guidelines for Telemetry Data Subscriptions over gNMI.]

Table 7: MX301 Feature Support (*Continued*)

Feature	Description
Layer 2 features	• Support for Layer 2 (L2) features, including: • Layer 2 protocol tunneling (L2PT) • Virtual private wire service (VPWS) • EVPN-MPLS and EVPN-VXLAN overlay networks • EVPN routing policies • Multiple VLAN Registration Protocol (MVRP) and Ethernet ring protection switching (ERPS) • VPLS over transport class tunnels • Headend termination of pseudowire services in a VPLS-enabled virtual switch • Q-in-Q tunneling • Basic DHCP snooping [See Layer 2 Protocol Tunneling, Overview of Headend Termination for EVPN VPWS, Multicast Support in EVPN-VXLAN Overlay Networks, Routing policies for EVPN, Understanding Multiple VLAN registration protocol (MVRP) for Dynamic VLAN Registration, Ethernet Ring Protection Switching Overview, Introduction to VPLS, Layer 2 Services on Pseudowire Service Interface Overview, Configuring Q-in-Q Tunneling, and DHCP Snooping.]

Table 7: MX301 Feature Support (*Continued*)

Feature	Description
Layer 3 features	• Support for Layer 3 (L3) features, including: • Network slicing with routing-based and firewall-based slice steering • Symmetric consistent hashing for symmetric load-balancing with correlated source IP and destination IP load-balancing hash-key (both forward and reverse direction) • OpenConfig-compliant Segment Routing-Traffic Engineering (SR-TE) sensor compliance • Adaptive load balancing (ALB) on aggregate Ethernet (AE) interfaces • Consistent hashing on flex-algo routes in a BGP network • Transit statistics for pseudowire and PWHT interfaces • GRE tunnels over PWHT interfaces • Seamless MPLS L3 support with pseudowire subscriber over redundant logical tunnel and pseudowire subscriber over logical tunnel. • Multicast routing protocols • MPLS Segment routing • Layer 3 (L3) unicast forwarding and GRE tunneling • RFC-2544-based benchmarking tests [See Hierarchical Class of Service for Network Slicing, Actions in Routing Policy Terms, YANG Data Model Explorer, Understanding Aggregated Ethernet Load Balancing, show interfaces extensive, MPLS Pseudowire Subscriber Logical Interfaces, MPLS Overview, Multicast Routing Protocols, Generic Routing Encapsulation (GRE), and Understanding RFC 2544-Based Benchmarking Tests.]

Table 7: MX301 Feature Support (*Continued*)

Feature	Description
MACsec	Support for MACsec bounded delay protection. Enable Media Access Control Security (MACsec) bounded delay protection to protect your network against man-in-the-middle attacks. When enabled, the device ensures that a frame is not sent after a delay of two seconds or more. MACsec periodically compares the number of frames transmitted to the number of frames received. If a frame is sent but not received within two seconds, such as during a man-in-the-middle-attack, the packet is dropped. [See bounded-delay.]
Network management and monitoring	Support for port mirroring. Analyzer is currently not supported. [See Port Mirroring and Analyzers.]
Protection against DDoS attacks	Support for distributed denial-of-service (DDoS) protection. DDoS protection is enabled by default. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]
Routing protocols	Support for L3 routing protocols. [See Routing Protocol Process Overview.]
Routing policy and firewall filters	Support for forwarding firewalls and 1GbE firewall support. Fabric loopback will be replaced by WAN stream loopback, which can impact mainline throughput. [See Understanding Firewall Filter Match Conditions and Service Filter Overview.]

Table 7: MX301 Feature Support *(Continued)*

Feature	Description
Security services	• Junos OS supports IKE to define mechanisms for key generation and exchange, and manage security associations (SAs). Inline IPsec is supported in Unified Services Framework (USF) mode. The IPsec architecture provides a security suite for the IPv4 and IPv6 network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. ESP in tunnel mode is supported. [See IPsec Overview and Enabling and Disabling Next Gen Services.]

Table 7: MX301 Feature Support (*Continued*)

Feature	Description
Services applications	• Support for inline services, including: • Inline NAT—NAT44 and NPTv6 • Inline softwires—Mapping of Address and Port with Encapsulation (MAP-E) • Mapping of Address and Port with Translation (MAP-T) • IPv6 rapid deployment (6rd) [See Inline NAT, Configuring Mapping of Address and Port with Encapsulation (MAP-E), and Configuring Inline 6rd.] • FlowTapLite support. [See Understanding the FlowTap and FlowTapLite Services.] • Inline video monitoring support, including: • Inline video monitoring using media delivery indexing (MDI) • Inline video monitoring of IPv4-over-MPLS flows to measure media delivery index (MDI) metrics • Inline video monitoring of IPv6 flows and IPv6-over-MPLS flows to measure media delivery index (MDI) metrics [See Understanding Inline Video Monitoring on MX Series Routers, video-monitoring, and show services video-monitoring mdi flows fpc-slot.] • Inline monitoring services support to provide the flexibility to monitor different streams of traffic at different sampling rates on the same interface. [See Inline Monitoring Services Configuration.] • Inline active flow monitoring support, including: • EVPN traffic

Table 7: MX301 Feature Support (*Continued*)

Feature	Description
	• MPLS, MPLS-IPv4, and MPLS-IPv6 traffic • IPv4 or IPv6 traffic, including traffic on next-hop based GRE tunnels and ps interfaces Both IPFIX and V9 formats are supported. [See Understand Inline Active Flow Monitoring.]
Subscriber management and services	• Support for subscriber services in Broadband Edge (BBE) and Broadband Network Gateway (BNG) solutions. Inline carrier-grade NAT (CGNAT) is currently not supported. [See Broadband Subscriber Services.] • Support for broadband network gateway (BNG) control and user plane separation (CUPS) architecture. Juniper BNG CUPS is a cloud application that separates the BNG control plane components from the native Junos OS. [See Juniper BNG CUPS Overview.] • Support for 5G Access Gateway Function (AGF). [See Access Gateway Function.] • Support for 5G Wireless CUPS user plane function (UPF). The UPF, supported on MX Series routers, handles all user data traffic according to the CUPS model. Junos Multi-Access UPF supports MX Series routers functioning as UPFs and provides high-throughput 5G fixed wireless and mobile wireless service in non-standalone (NSA) mode. [See MX Series Router As Junos Multi-Access User Plane.]

Table 7: MX301 Feature Support (*Continued*)

Feature	Description
Software installation and upgrade	- Secure boot and common BIOS support. The Secure Boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms. [See Junos OS Overview.] - Firmware upgrade support. [See request system firmware upgrade.] - Support for secure zero-touch provisioning (SZTP). Use RFC 8572-based SZTP to bootstrap your remotely located network devices that are in a factory-default state. SZTP enables mutual authentication between the bootstrap server and the network device before the remote network device is accessed for initiating zero-touch provisioning. [See Secure Zero Touch Provisioning.]

Chassis

- Chassis management support (MX301)**—The MX301 router supports the following chassis management features:
 - 1+1 redundancy mode for the power supply units (PSUs)
 - Environment monitoring policy for sensor components on Routing Engine, CPU, memory, fan, and line cards
 - Humidity sensor to display relative humidity and temperature using the `show chassis routing-engine humidity` command

[See [show chassis routing-engine humidity](#).]

EVPN

- Weighted ECMP for EVPN-MPLS Type 5 routes (MX301, MX304, MX10004, and MX10008)**—Use weighted equal-cost multipath (WECMP) for EVPN-MPLS Type 5 routes by advertising generalized weights from gateway counts with the EVPN Link Bandwidth extended community. This

advertisement distributes traffic across next hops and prevents congestion when leaf capacities differ. To configure an export policy, use `set policy-options policy-statement policy-name` then `aggregate-weight [multiplier one through 32]` and `set routing-instances instance-name protocols evpn ip-prefix-routes export policy-name`. Use `show evpn ip-prefix-database` extensive to verify weights, balance factors, and feature state. If weights are missing or inconsistent from the Type 5 route of any leaf device, the border leaf device generates a debug log and follows regular ECMP forwarding to the leaf devices.

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

- **EVPN-VXLAN Type-5 support over IPv6 underlay with MAC-VRFs (MX304 and MX960)**—You can deploy EVPN-VXLAN over an IPv6 underlay and advertise EVPN Type-5 IP-prefix routes to extend L3 connectivity. Use MAC-VRF instances to provide L2 services while distributing L3 prefixes end-to-end across an IPv6-only fabric, improving scalability and aligning the network with dual-stack or IPv6-centric designs. To integrate L2 and L3 services, configure IPv6 underlay routing, create EVPN MAC-VRF instances, and enable Type-5 route advertisement. This capability applies to MAC-VRF instances only. Additionally, you can configure IPv6 transport to enable VXLAN EVPN Type-5 route advertisement and SRv6 stitching for seamless interworking between EVPN-VXLAN fabrics and SRv6 transport domains. This capability allows you to build IPv6-only data-center and WAN fabrics that deliver scalable Layer-2 services using MAC-VRF instances and flexible Layer-3 service extension through Type-5 routes.

[See [EVPN-VXLAN with an IPv6 Underlay](#).]

High Availability

- **Support for single-hop BFD echo-lite sessions in distributed and inline mode (MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—You can configure single-hop BFD echo-lite sessions in inline and distributed modes to monitor forwarding path availability for directly connected next hops. Unlike regular echo sessions, echo-lite sessions work even if the next-hop device is not configured for BFD. Use the following configuration command to configure BFD echo-lite sessions:

```
set routing-options static route address bfd-liveness-detection echo-lite minimum-interval interval
```

[See [Understanding BFD Echo and Echo-Lite Modes](#).]

- **Unified ISSU with enhanced mode for Layer 3 on MX10K-LC4800 and MX10K-LC4802 line cards (MX10004 and MX10008)**—You can perform unified ISSU with enhanced mode for Layer 3 on MX10K-LC4800 and MX10K-LC4802 line cards. Enhanced mode is an ISSU option that prevents packet loss during the unified ISSU process.

To use unified ISSU with enhanced mode, enter the following command:

```
request system software in-service-upgrade image reboot verbose enhanced-mode
```

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

Interfaces

- **QSFP-100G coherent ZR optics support (MX204)**—The optics enhancements include application selection and configuration of target output power. You can view the advertised applications and switch between the applications.
- **QSFP-100G coherent ZR optics support (MX240, MX480, MX960 with MPC10E-10C-MRATE line card)**—Manage optical transport links efficiently with QSFP-100G coherent ZR optics. QSFP-100G-ZR optics do not support the 'none' forward error correction (FEC) mode.
- **QSFP-100G-LRBD coherent ZR optics support (MX304)**. Manage optical transport links efficiently with QSFP-100G-LRBD coherent ZR optics.
- **QSFP-100G-ERBD coherent ZR optics support (MX304)**. Manage optical transport links efficiently with QSFP-100G-ERBD coherent ZR optics.
- **Interface, Port speeds, and channelized subports (MX301)**—You can manage PIC ports with port profiles that select active ports and their speeds. This feature provides precise control of utilization and performance. Configure profiles at the PIC or per-port level using `pic-mode` and `speed` options. Supported rates include 1, 10, 25, 40, 50, 100, 200, and 400 GE. For channelized ports, specify interfaces per physical port with the CLI: `set chassis fpc 0 pic 0 port <> speed <> number-of-subports <>`. This configuration helps you tailor capacity and interface density to your deployment.

Junos OS API and Scripting

- **Support for libslax 3.1.6 and SLAX version 1.3 (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—We've upgraded the libslax library to version 3.1.6, which corresponds to SLAX version 1.3. The upgraded library includes:
 - Slax processor enhancements including a new mode, additional options, and simplified argument parsing
 - New libslax extension library functions

- Improved SLAX syntax options
- New SLAX functions and enhancements to existing functions and statements
- Hexadecimal numbers support in SLAX scripts

This update aligns SLAX processing with contemporary scripting conventions, ensuring efficient, readable scripting while maintaining backward compatibility.

[See [libslax Distribution Overview](#).]

Precision Time Protocol (PTP)

- **Memory logging infrastructure for enhanced timing application debugging (MX240, MX480, MX960, MX2010, and MX2020)**—Enhance your network's debugging capabilities with the Memory Logging Infrastructure for timing applications. This feature enables you to log and categorize critical and operational events, improving system resilience and troubleshooting efficiency. You can manage logs through specific CLI commands and configure options like dump-on-wrap to maintain log records during system restarts. By utilizing this infrastructure, you streamline memory logging across platforms, ensuring comprehensive performance management and operational visibility, particularly with timing processes such as `clksyncd` and `clksyncm`.

[See [Understanding Memory Logging Infrastructure for Timing Applications](#).]

- **Phase adjustment mitigation for G.8275.1 hybrid mode (MX240, MX480, MX960, MX2010, and MX2020)**—Use the phase adjustment mitigation feature to protect downstream timing from GNSS-induced phase jumps in hybrid deployments. This feature prevents phase adjustments to downstream Precision Time Protocol (PTP) nodes caused by a significant phase jump of the upstream T-GM or T-BC PTP node. Use the protocols `ptp phase-error-limit relative-threshold <nanoseconds>` configuration statement to ignore phase adjustments that exceed the threshold, with SNMP traps and syslogs for visibility. Optionally, use protocols `ptp phase-error-limit force-synce-holdover` and protocols `ptp phase-error-limit max-holdover-time <minutes>` to place PTP and Synchronous Ethernet in holdover and control duration. The system applies a ten-measurement-window grace period before resuming adjustments. Monitor status with `show ptp phase-error-monitoring status` and `show ptp global-information` commands. This control applies to PTP-originated adjustments.

[See [Phase Adjustment Mitigation in Hybrid Mode](#).]

- **PTP passive port monitoring with SNMP MIB support and timeReceiver performance monitoring (MX304)**—You can enhance router performance monitoring through Precision Time Protocol (PTP) passive port monitoring and SNMP MIB integration. Use `show ptp global-information` and `show ptp passive-port-monitor-status` commands to display performance statistics. Passive monitoring alerts you to potential fiber asymmetries or clock failures, improving network reliability by identifying issues when phase thresholds are exceeded.

TimeReceiver performance monitoring feature allows you to collect detailed metrics for active PTP timeReceiver ports using IEEE 1588-2019 standards, including timeTransmitter–timeReceiver delay and synchronize message transmission. Use the `show ptp performance-monitor status` command to view the status of timeReceiver port performance monitoring.

[See [PTP Passive Port Performance Monitoring](#) and [PTP TimeReceiver Port Performance Monitoring](#).]

- **Support for PTP G.8275.2 enhanced profiles for Telecom phase and time synchronization (MX10008 and MX10004 with JNP10K-LC9600 line cards)**—You can enable Precision Time Protocol (PTP) G.8275.2 enhanced profiles to achieve compliance with ITU-T G.8273.4 standards for precise phase and time synchronization in Telecom applications. This feature supports PTP over both IPv6 and IPv4 unicast, ordinary and boundary clocks, and unicast negotiation, facilitating accurate synchronization over wide area networks. To configure this enhanced profile, enable the `g.8275.2.enh` statement in the `[edit protocols ptp profile-type]` CLI hierarchy.

[See [G.8275.2 Enhanced Profile](#).]

- **Support for PTP over IRB for G.8275.2 profile (MX10004 and MX10008 with JNP10K-LC480 line cards)**—Precision Time Protocol (PTP) over Integrated Routing and Bridging (IRB) for G.8275.2 profile enables precise timing synchronization in Remote PHY deployments by allowing PTP timeTransmitter configuration on IRB interfaces. This feature supports clustered remote PHY devices (RPD) environments with a single PTP timeTransmitter serving multiple timeReceivers, aggregated Ethernet integration within bridge domains, and redundancy through primary/secondary link configuration. It ensures telecom-grade synchronization for Data-over-Cable Service Interface Specification (DOCSIS) and emerging services like wireless backhaul.

[See [PTP over IRB for G.8275.2 Enhanced Profile](#).]

- **Support for Synchronous Ethernet, Synchronous Ethernet over LAG, G.8275.1, and G.8275.1 over LAG (MX10004 and MX10008 routers with JNP10K-LC4802 line card)**—Use Synchronous Ethernet and G.8275.1 over Link Aggregation Groups (LAG) on MX10004 and MX10008 routers equipped with JNP10K-LC4802 line cards to enhance network performance, stability and redundancy. Synchronous Ethernet provides precise frequency synchronization over Ethernet networks for time-sensitive applications, ensuring stable and reliable communication. Link Aggregation Group (LAG) increases bandwidth and provides redundancy by combining multiple physical links into a single logical connection. This enhances network performance and ensures high availability in case one link fails.

The G.8275.1 profile, based on ITU-T G.8275, enables the distribution of phase and time with full timing support. Operate all devices in combined or hybrid mode, where both Precision Time Protocol (PTP) and Synchronous Ethernet are enabled.

[See, [Timing Features with JNP10K-LC480 and JNP10K-LC4800 Line Cards on MX10004 and MX10008](#).]

Junos Telemetry

- **Support to monitor soft-GRE tunnels and sessions (MX240, MX480, and MX960)**—Use Junos telemetry sensors to monitor soft-GRE tunnel or session information. Juniper's YANG modules for native GRE sensors are available at Juniper's telemetry [GitHub](#) repository. To configure data collection from the sensors, use the Junos CLI to provision sensors for collecting specific data. Subscribe to the sensor information you want to collect; data is streamed over gRPC or UDP connections. Use the protocol buffer's files to decode the streamed data on the collector. You can retrieve the following soft-GRE sensor information:
 - At the system level, you can obtain the total count of GRE static or dynamic tunnels. You can also obtain information on session types (static, PPPoE, DHCP, and VLAN).
 - At the tunnel level, you can retrieve detailed information about interface name, address types, local or remote addresses, routing instance, and session counts.
 - You can retrieve information about traffic data for inbound and outbound packets, bytes, and rates (packets or bits per second) for each tunnel.

This data helps monitor and manage the performance of GRE tunnels and their associated sessions.

The following resource paths are supported:

- /state/services/soft-gre/counters/static-tunnels
- /state/services/soft-gre/counters/dynamic-tunnels
- /state/services/soft-gre/counters/static-sessions
- /state/services/soft-gre/counters/pppoe-sessions
- /state/services/soft-gre/counters/dhcp-sessions
- /state/services/soft-gre/counters/vlan-sessions
- /state/services/soft-gre/counters/user-sessions
- /state/services/soft-gre/logical-interfaces/logical-interface/
- /state/services/soft-gre/logical-interfaces/logical-interface/name
- /state/services/soft-gre/logical-interfaces/logical-interface/static-tunnels
- /state/services/soft-gre/logical-interfaces/logical-interface/dynamic-tunnels
- /state/services/soft-gre/logical-interfaces/logical-interface/static-sessions
- /state/services/soft-gre/logical-interfaces/logical-interface/pppoe-sessions

- /state/services/soft-gre/logical-interfaces/logical-interface/dhcp-sessions
- /state/services/soft-gre/logical-interfaces/logical-interface/vlan-sessions
- /state/services/soft-gre/logical-interfaces/logical-interface/user-sessions
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/name
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/local-address
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/ remote-address
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/routing-instance
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/static-sessions
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/pppoe-sessions
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/dhcp-sessions
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/vlan-sessions
- /state/services/soft-gre/tunnel-interfaces/tunnel-interfaces/user-sessions
- /state/services/soft-gre/session-interfaces/session-interface/
- /state/services/soft-gre/session-interfaces/session-interface/name

For more information, see [Sensors and Sensor Paths, Guidelines for Streaming Telemetry Data Over UDP | Junos OS | Juniper Networks](#), [Establish a Dial-out Telemetry Connection | Junos OS | Juniper Networks](#), [Decoding Junos Telemetry Interface Data With UNIX Utilities](#), and [Junos YANG Data Model Explorer](#).

- **Sensor path support for the table connection disable-metric-propagation leaf (ACX710, MX10008, and MX10016)**—By default, the OpenConfig protocol sets the destination protocol metric based on the source protocol metric. Set the disable-metric-propagation leaf to true to stop this behavior. The device then sets the metric to zero or a policy-defined value. Subscribe to the following resource paths to obtain sensor-specific information:
 - /network-instances/network-instance/table-connections/table-connection/config/disable-metric-propagation
 - /network-instances/network-instance/table-connections/table-connection/state/disable-metric-propagation

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

- **OpenConfig IS-IS enhancements including graceful restart and maximum ECMP paths (MX204, MX240, MX301, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008,**

MX2010, MX2020)—Configure the missing leaves for graceful restart and maximum ECMP paths to achieve compliance with OpenConfig data models for IS-IS protocols. This step ensures remote management across different vendors without requiring device login, by using specific CLI commands to set configurations that include OpenConfig network instances and graceful restart options.

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

- **Stream telemetry data in gNMI-based message format over UDP (MX204, MX240, MX301, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, and QFX5210)**—Junos OS uses a dial-out mechanism to send telemetry data to a collector over UDP. The message format is defined in the `jnx_gnmi_over_udp.proto` file. This mechanism supports only STREAM mode with SAMPLE as subscription mode. The message contains full key name and value pair information so the collector does not require data models for processing or consuming the telemetry data.

No Link Title, No Link Title, No Link Title, No Link Title, No Link Title, and [Junos YANG Data Model Explorer](#).]

MACsec

- **Automatic adjustment of MTU for MACsec overhead (MX304, MX960, MX2020, MX10003, and MX10008)**—Use this feature to automatically adjust the maximum transmission unit (MTU) for the Media Access Control Security (MACsec) overhead. Without this feature, you must adjust the interface MTU and the protocol MTU manually.

Use this feature to ensure the interface or protocol MTU is adjusted properly to account for the MACsec overhead. This feature is disabled by default. To enable this feature, first enable MACsec. Then configure the `enable-auto-mtu-update` statement at the `[edit security macsec]` hierarchy level. This feature applies to physical interfaces, logical interfaces, and physical interfaces that are members of aggregated Ethernet interfaces.

[See [Media MTU and Protocol MTU](#).]

MPLS

- **Pseudowire over Aggregated Ethernet (MX304)**—Terminate pseudowire connections directly on the aggregated Ethernet (AE) interface to bypass the downstream loopback penalty associated with pseudowire headend termination (PWHT). Pseudowire over AE (PSoverAE) supports the same features as PWHT and pseudowire over RLT (PSoRLT). Some features such as class of service (CoS) are more streamlined compared with PWHT and PSoRLT.

[See [MPLS Pseudowires Configuration](#) and [Aggregated Ethernet Configuration](#)].

- **BGP VPN Option C MPLS-over-SRv6 (MX204, MX240, MX301, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Use this feature to

integrate SRv6 with MPLS-based BGP Layer 3 VPN services. Use SRv6 SIDs for routing in SRv6 domains while preserving MPLS label operations in MPLS domains. This feature enables MPLS-based BGP Layer 3 VPN services to interwork seamlessly across SRv6 by leveraging BGP label unicast (BGP-LU) for IPv4 loopback reachability and signaling SRv6 SIDs alongside MPLS labels using the Prefix-SID attribute. Configure static or dynamic end-dtm-SIDs and enable BGP-LU to advertise or accept SRv6 services. Manage ingress, transit, and egress routes, including label push, swap, and encapsulation. This configuration enhances scalability across multi-domain networks.

[See [Understand BGP VPN Option C MPLS over SRv6.](#)]

- **Enhancements to MPLS auto-bandwidth (MX204, MX240, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—Use dynamic bandwidth allocation for RSVP label-switched paths based on traffic volume, providing granular control and flexibility in managing network resources. The `polling-profile` statement under `[edit protocols mpls statistics auto-bandwidth]` supports multiple polling profiles with different intervals tailored to specific label-switched paths. Use this configuration to optimize bandwidth adjustments to meet diverse service requirements. Apply `adjust-threshold` settings for precise overflow and underflow triggers. Enhanced CLI commands define polling profiles, associate them with LSPs, and configure thresholds, improving monitoring and troubleshooting in complex network environments.

[See [Using Polling Profiles for Automatic Bandwidth Allocation for RSVP LSPs.](#)]

Network Management and Monitoring

- **Ephemeral database default commit synchronize model changed to synchronous (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—We've changed the default commit synchronize model for the ephemeral database from the asynchronous model to the synchronous model. The synchronous model ensures better synchronization of ephemeral configurations across Routing Engines or Virtual Chassis members by processing commits synchronously. With this change, only the synchronous model supports synchronizing ephemeral data on devices that have graceful Routing Engine switchover (GRES) or nonstop active routing (NSR) enabled.

[See [Understanding Ephemeral Database Commit Synchronize Models.](#)]

NextGen Port Extender (NGPE)

- **NextGen Port Extender (MX10004, MX10008, and MX304)**—The NextGen Port Extender is a solution for managing an aggregation network environment, in which a central aggregation device connects to multiple satellite devices. The interfaces on each satellite device become extensions of the aggregation device and the whole cluster appears, from a management perspective, as a single device. EVPN-VXLAN tunneling transports all traffic between aggregation and satellite devices. The MX10004 and MX10008 are supported only with MX10K-LC4800 or MX10K-LC9600 FPCs.

[See <https://www.juniper.net/documentation/us/en/software/junos/ngpe/index.html>.]

- **VPLS support for NextGen Port Extender (MX304, MX10004, and MX10008)**—Use the NextGen Port Extender to provision VPLS services on the extended ports of the satellite device. The MX304, MX10004, and MX10008 routers act as the aggregation device.

[See [Introduction to Configuring VPLS](#).]

Post-Quantum Cryptography (PQC)

- **PQC signatures for software images with off-box verification (ACX Series, EX Series, MX Series, QFX Series, and SRX Series)**—Use post-quantum cryptography (PQC) signatures to ensure software images remain unaltered, protecting integrity and guarding against quantum threats. The images comply with algorithms recommended by Commercial National Security Algorithm 2.0 (CNSA 2.0):
 - ML-DSA-87 PQC algorithm for digital signatures
 - SHA-512 for hashing

For off-box verification, request the PQC signature from the support portal. Confirm the image's SHA-512 hash, retrieve the public key from the certificate, and validate the signature with your chosen verifier. PQC signatures provide additional security beyond existing legacy signatures.

[See [PQC Signatures for Software Images](#).]

Routing Protocols

- **AS loop check in BGP Nnetworks (MX304, MX10004, and MX10008)**—We have enabled Autonomous systems (AS) path loop check for external BGP (EBGP) and internal BGP (IBGP) sessions by default. The loop check is made in the BGP peer's AS path domain. Use this feature to configure and manage AS path loop detection on Junos devices.

You can disable AS path loop check for IBGP sessions including all routing instances using the `no-loop-check` statement at the `[edit protocols bgp defaults ibgp]` hierarchy level.

[See [no-loop-check](#).]

- **BGP prefix limit based on route target to limit VPN prefixes (MX304, MX10004, and MX10008)**—In typical L3VPN deployments, you can limit routes at the customer edge peer level with the `prefix-limit` configuration for a BGP peer family. We have shifted this control to a central location such as the route reflector or an AS boundary router (ASBR) so that routes originating at all sites in a VPN are taken into account. BGP maintains and enforces the prefix limit as specified by the route target communities originating at various VPN sites to limit the number of prefixes a BGP peer can advertise or receive to conserve resources.

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

- **Configure BGP keepalive value independent of holdtimer value (MX2008)**—BGP uses hold time to terminate unresponsive sessions. The hold time resets each time a BGP message is received on a BGP connection. If messages stop, a keepalive timer triggers a keepalive message. By default, this keepalive timer is one third of the negotiated BGP hold time, if greater than zero. You can now configure the keepalive timer independently of the hold time between 1 second and 21845 seconds. Include the `keepalive-time` statement at the `[edit protocols bgp]` hierarchy level.

[See [keepalive-time](#).]

- **Generate static RT-Constrain route based on community or wildcard (MX304, MX10004, and MX10008)**—When the RT-Constrain feature is partially deployed in a network, the resource saving benefit is lost. We have extended the static RT-Constrain feature to generate host static RT-Constrain entries from fully qualified route targets configured in the routing policy. You can assign BGP communities or a wildcard route target on the static RT-Constrain route. You can also configure the static RT-Constrain route's origin AS in the network layer reachability information (NLRI) while retaining the global AS number.
- **IS-IS multi-instance support on a single interface (ACX5448, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX10002-60C, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016)**—We have enhanced the IS-IS multi-instance feature to support multiple IS-IS instances on the same logical interface with instance identifier TLV 7.

Include the `instance-id` statement at the `[edit protocols isis-instance name]` hierarchy level.

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

- **Path quality profile sharing for GLB multi-link support on IP Fabric (MX10004)**—While we support larger Clos networks and more GPUs, the TH5 chipset can only support 64 profiles. In Clos networks with five or more stages, some nodes, like super spines exceed 64 next-next-hop nodes. We can reuse the profiles under specific conditions to support more than 64 next-next-hop nodes. We support GLB for Clos networks with profile sharing in hyperscaler artificial intelligence/machine learning (AI/ML) fabrics containing tens of thousands of leaves or GPUs.

[See [profile-sharing](#).]

- **Replace BGP AS path to maintain network interoperability (MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—Define a routing policy to match and replace a list of autonomous systems (AS) numbers or private AS numbers with the local AS number of the BGP peering session to maintain network interoperability. This configuration works only on AS sequences and not on AS sets. In route-reflector scenarios, in addition to using external BGP (EBGP), enable internal BGP (IBGP) to leverage this capability. Include the policy action `as-path-replace as-list | private` statement at the `[edit policy-options policy-statement statement-name then]` hierarchy level to activate the feature.

[See [Autonomous Systems for BGP Sessions](#).]

- **Selectively disable NH validation based on community, route target, or RD (MX304, MX10004, and MX10008)**—Define a BGP import policy to selectively disable next hop resolution for routes that match a specified community, route target, or route distinguisher (RD). The policy sets the next hop to fictitious instead of an indirect next hop address, which avoids resolving the next hop for routes that match the community specified in the policy.

Services Applications

- **RFC 8402 SR support for TWAMP probes (MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—We have added support in Two-Way Active Measurement Protocol (TWAMP) for segment routing (SR) as defined in RFC 8402, which broadly specifies the SR architecture. We support two types of SR for TWAMP probes:
 - **SR-MPLS:** Uses a list of labels, each representing a segment end node.
 - **SRv6:** Uses a type 4 IPv6 routing header introduced in RFC 8754, with each segment end node represented as an IPv6 address or IPv6 segment identifier (SID).

You can specify the list of SR-MPLS or SRv6 segments for a TWAMP probe to reach the reflector and from the reflector to the client. This return path information is embedded in the probe itself by using the extensions proposed in *Simple TWAMP (STAMP) Extensions for Segment Routing Networks*, draft-ietf-ippm-stamp-srpm, namely, the return path TLV and its return segment list sub-TLVs. The information is embedded depending on the segment routing type. The device timestamps the TWAMP probes in either the Routing Engine or the Packet Forwarding Engine.

To configure this feature, include the source-routing statement at the `[edit services rpm twamp client control-connection name test-session session-name]` hierarchy level.

[See [Understand Two-Way Active Measurement Protocol](#) and [source-routing](#).]

- **Use FQDN targets in RPM probes (MX204, MX240, MX304, MX480, and MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—Configure a real-time performance monitoring (RPM) probe test with a fully qualified domain name (FQDN) target (instead of an IPv4 or IPv6 address) to adapt to endpoint address changes. The device resolves the FQDN through the Domain Name System (DNS) at the start of each test and then sends probes to the

resolved address at each interval. Telemetry first displays the FQDN until resolution and then shows the destination address to help you track changes. If DNS resolution fails, the probe results retain the FQDN as the target address. Configure the `fqdn` or the `fqdnv6` option at the `[edit services rpm probe owner test test-name target]` hierarchy level.

[See [target](#).]

- **Traffic selector support for inline IPsec (MX304, MX10004, and MX10008)**— Inline IPsec encrypts and decrypts IPsec tunnel traffic directly within the Packet Forwarding Engine, by eliminating the need for a dedicated services card. A traffic selector is an agreement between IKE peers to permit traffic through a tunnel if the traffic matches a specified pair of local and remote addresses.

[See [Inline IPsec](#).]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **CBF for SRv6-TE (MX10003, MX10008, and MX10016)**—We extend the existing color-based forwarding (CBF) functionality to Segment Routing for IPv6—Traffic Engineering (SRv6-TE) enabling you to adapt to complex network demands. Use this feature to steer traffic across multiple transport tunnels based on class of service (CoS). This approach improves route selection, next-hop resolution, and service quality. Resolver enhancements support SRv6 Segment Identifiers (SIDs) across multipath routes. Use the `preserve-next-hop-hierarchy` configuration statement to prevent misrouting.

[See [preserve-next-hop-hierarchy \(SRv6-TE\)](#).]

- **Performance measurement for SR-MPLS and SRv6 paths (MX204, MX240, MX301, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, and MX2020)**—You can now enhance network path selection using the Two-Way Active Measurement Protocol (TWAMP) for performance-based metrics such as latency on Segment Routing over MPLS (SR-MPLS) and Segment Routing over IPv6 (SRv6) paths. This feature allows you to enable end-to-end delay measurement for Segment Routing paths with static segment lists. Employ the new CLI options to configure probe intervals, counts, and advertisement intervals. This helps monitor network performance efficiently, allowing you to make informed path-selection decisions and maintain optimal service levels, particularly in performance-critical environments such as financial networks.

[See [delay measurement](#) and [show spring-traffic-engineering performance-measurement](#).]

- **Support for UHP in IS-IS SR-MPLS (MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—Use ultimate-hop popping (UHP) with IS-IS or OSPF so the egress provider edge can process its own node SID. IS-IS advertises a node SID with the P flag set and E flag unset. In controller-driven segment routing—traffic engineering (SR-TE), the controller inserts the egress provider edge node SID beneath the SR-TE binding SID (BSID). If the BSID route fails on the penultimate hop, the egress provider edge might see its own node SID as the top label instead of penultimate-hop-popping (PHP). With the P flag set, the provider edge expects UHP and processes its MPLS label. Include the `ultimate-hop-popping` statement at the `[edit protocols isis source-packet-routing]` hierarchy level.

[See [ultimate-hop-popping](#)ultimate-hop-popping.]

- **SRv6-TE route resolution over BGP without IGP (MX10003, MX10008, and MX10016)**—A Segment Routing for IPv6 (SRv6) tunnel consists of paths with Segment Identifiers (SIDs) over an interior gateway protocol (IGP) that steer traffic to a traffic-engineering (TE) path. If IGP is not available, you can configure these SIDs statically and advertise them through external BGP (EBGP). We support this feature on both classic and micro SRv6 SIDs.

Networks that deploy a network orchestrator to steer transit traffic onto a TE path and advertise these transit prefixes using BGP color community typically don't have a service SID. In this case, the last SID must not be removed. The ingress SRv6 TE tunnel acts as a transit tunnel to forward transit traffic with SRv6 encapsulation.

Include the `no-remove-srv6-last-sid` statement at the `[edit protocols source-packet-routing]` hierarchy level and the `use-ingress-routes-as-transit` statement at the `[edit protocols source-packet-routing srv6]` hierarchy level.

[See [no-remove-srv6-last-sid](#).]

- **Delay normalization for OSPF Flexible Algorithm metrics and advertisements across IGP instances (MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—Use delay normalization to compute and advertise a normalized delay metric for Flexible Algorithm, to improve path-selection consistency across all interior gateway protocol (IGP) instances. The device normalizes each received delay, compares each value with the previously saved normalized value, and triggers link-state advertisement (LSA) generation when the values differ.

Delay normalization is disabled by default. To enable and configure delay normalization, use the `normalize interval offset` statement at the `[edit protocols ospf area interface delay-measurement]` hierarchy level.

[See [delay-measurement \(Protocols OSPF\)](#) and [How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering](#).]

- **Selectively control per-prefix backup paths with OSPF import policy (MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—You can selectively enable backup paths for specific prefixes to optimize redundancy and resource utilization. By default, a configured backup path applies to all prefixes. To exclude specific prefixes or ranges, create an OSPF import policy and configure the `no-backup` option in the `then` clause of the policy to suppress backup path installation for matching routes. You can reserve backup protection for critical prefixes while preventing unnecessary backups for others.

[See [Understanding Backup Selection Policy for OSPF Protocol](#).]

- **Preference-based path selection for L-OSPF Flexible Algorithm routes (MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and**

MX10016)—You can control path selection by configuring the preference for L-OSPF Flexible Algorithm routes in `inetcolor.0` and `mpls.0`.

Configure `flex-algorithm-preference` statement at the `[edit protocols ospf]` hierarchy level to prioritize desired routes and improve traffic engineering across IP and MPLS domains.

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

- **Policy-based redistribution of OSPF prefix SIDs across IGP instances (MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)**—You can redistribute Segment Routing (SR) prefix-SIDs across OSPF IGP instances using route policy without explicitly specifying a prefix-segment index. This feature standardizes SR labels across instances and improves operational efficiency. Configure a policy with the `from prefix-segment` statement to match routes carrying prefix-segment information. In the `then` clause, use `prefix-segment redistribute` to inherit segment information from the matched route. We also support stitching `mpls.0` routes to enable interoperability between different IGP instances.
- **Static SID configuration in SRv6 Manager (MX10003, MX10004, and MX10016)**—Configure SRv6 classic, micro node, adjacency SIDs, along with classic END and END-X SIDs, and install them in the routing table without using interior gateway protocols (IGPs) such as IS-IS. Advertise these static routes through a BGP export policy for path computation. This configuration enables controllers to receive static node and adjacency SIDs over BGP and compute SR-TE paths across a domain that does not use IS-IS.

[See [Understanding SRv6 Static Segment Identifier](#).]

Subscriber Management and Services

- **BGP support over dynamic PPPoE subscriber interfaces (MX480 and MX960)**—You can enable BGP over dynamic PPPoE interfaces for IPv6 family on routers with MPC2, MPC5, and MPC7 line cards. To enable this support, configure routing services with the `routing-service` statement in both the PPPoE subscriber dynamic profile and the dynamic profile for the underlying VLAN interface. You can also enable routing services for subscribers over pseudowire interfaces when a non-demux dynamic underlying VLAN is configured.

[See [Enabling BGP over Dynamic PPPoE Subscriber Interfaces](#) and [Example: Configuring IPoE and PPPoE Subscribers on the Same Access Interface with Routing Services Support](#).]

- **Operational continuity for pseudowire subscribers over RLT (MX304, MX960, and MX10004)**—You can instantiate pseudowire subscriber logical interfaces over a redundant logical tunnel (RLT) interface in active-active or active-backup mode to maintain operational continuity. This configuration ensures uninterrupted service by activating new subscribers and keeping existing subscribers operational if an RLT member interface goes down due to a Packet Forwarding Engine disablement. Subscribers remain up as long as at least one RLT member link is on an active Packet Forwarding Engine.

[See [Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview](#).]

- **MAC address validation support for IPv6 interfaces (MX204, MX240, MX480, MX960, MX10004, and MX10008)**—You can enable MAC address validation feature on interfaces configured with IPv6 addresses. This feature allows the router to validate that incoming packets contain both a trusted IP source and a MAC source address. You can configure the following two modes to take action based on validation:
 - Strict: Drops packets that fail validation.
 - Loose: Forwards packets even if validation fails.

You can enable MAC address validation on both static and dynamic interfaces. For static interfaces, use the `set interfaces interface-name unit logical-unit-number family inet6 mac-validate (loose | strict)` command. For dynamic interfaces, use the `set dynamic-profiles profile-name interfaces demux0 unit logical-unit-number family inet6 mac-validate (loose | strict)` command.

[See [MAC Address Validation for Subscriber Interfaces Overview](#), [Configuring MAC Address Validation for Subscriber Interfaces](#), [mac-validate](#), and [mac-validate \(Dynamic IP Demux Interface\)](#).]

- **Support for IPv6 subscribers over Soft-GRE tunnels (MX240, MX480, and MX960)**—Wi-Fi Access Gateway (WAG) supports IPv6 Soft Generic Routing Encapsulation (Soft-GRE) tunnel creation for carrying IPv4 and IPv6 traffic. The following IPv4 Soft-GRE tunnel features also support IPv6 Soft-GRE tunnels:
 - Static subscriber support over statically configured Soft-GRE tunnels
 - PPPoE support over Soft-GRE tunnels with active/active and active/backup RLT modes
 - Retrieve soft-GRE information through Junos telemetry

You can use the existing Soft-GRE tunnel configuration CLI commands for IPv6 tunnels.

[See [Wi-Fi Access Gateways](#), [Static Subscriber Support over Statically Configured Soft GRE Tunnels](#), [soft-gre](#), [show services soft-gre tunnel](#), [Junos Telemetry Interface User Guide](#), and [Junos YANG Data Model Explorer](#).]

Additional Features

We've extended support for the following features to these platforms.

- **On-device packet capture** (MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004). Junos devices support filtering and mirroring incoming and outgoing packets, sending those packets to the CPU, and saving them into a file. This feature, on-device packet capture, can help you with protocol and application analysis, debugging, troubleshooting, network forensics, audit trails, and network attack detection.

[See [On Device Packet Capture](#).]

- **Supported transceivers, optical interfaces, and DAC cables.** (EX4100, EX4400, EX4400-24X, EX4400-EM-4Y, EX4650, MX304, MX10004, MX10008, QFX5120, and SRX4700). Select your product in the [Hardware Compatibility Tool](#) 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 optical transceiver becomes available.
- **Support for performance monitoring and TCA** (MX240, MX480, and MX960). Use performance monitoring for MX240, MX480, and MX960 routers to measure current and historical metrics. The routers accumulate these metrics into 15-minute and 1-day intervals. You can configure the 15-minute interval length. You can view these metrics by using the [show interfaces transport pm](#) command. This approach helps you manage optical transport links more efficiently.
- **Two-Way Active Measurement Protocol (TWAMP) monitoring service (RFC 5357) hardware timestamp support to enable Flex Algo and SR-MPLS support** (MX Series platforms that support the MPC10E, MPC11E, JNP10K-LC4800, and the JNP10K-LC9600)

[See the offload-type inline-timestamping option of the [test-session](#) statement.]

- **400ZR and 400G OpenZR+ optics application selection enhancements** (MX304)

[See [400GbE ZR and 400GbE ZR+](#) and [Application Selection](#).]

- **Renaming OpenSSH implementation to JSSH (all platforms)**—The OpenSSH implementation in Junos OS is renamed "JSSH" to highlight its customized nature.

[See [ssh](#).]

What's Changed

IN THIS SECTION

- [Dynamic Host Configuration Protocol](#) | 85
- [Junos Telemetry](#) | 85
- [General Routing](#) | 85
- [MPLS](#) | 88
- [Network Management and Monitoring](#) | 88
- [Platform and Infrastructure](#) | 89
- [Routing Protocols](#) | 89
- [Subscriber Access Management](#) | 89

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

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

Junos Telemetry

- **Standards Compliance for Data Type**— To comply with the new nomenclature standards for YANG models, the sensor values with data type "Enumeration" that were previously in uppercase letters are now in lowercase letters.

Table 8: Data Type

Leaf	Resource Path	Data Type: Old Format	Data Type: New Format
poe-management	/state/poe/controllers/ controller/poe- management	Enumeration • STATIC • DYNAMIC • CLASS	Enumeration • static • dynamic • class

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

General Routing

- A new counter **Sessions hit due to high rate** is added to `show services service-sets screen-session-limit-counters` command for all subscriber traffic. This counter tracks the sessions that come up on the screen irrespective of the `alarm-without-drop` configuration. When "alarm-without-drop" option is disabled, all the counters display updated statistics. When `alarm-without-drop` is enabled, then: - The

screen-drop counters on "show services service-sets statistic screen-drop" command do not increase.

- The "sessions hit due to high rate" value is displayed.

[See [alarm-without-drop \(IDS Screen Next Gen Services\)](#), [show services service-sets statistic screen-drops \(Next Gen Services\)](#), and [show services service-sets statistic screen-session-limit-counters \(Next Gen Services\)](#).]

- **SFP Optics LOS alarms (MX Series)**—SFP Optics do not support Tx laser disabled alarm, Tx loss of signal functionality alarm, and Rx loss of signal alarm as diagnostics output.

[See [show interfaces diagnostics optics](#).]

- When you run the request vmhost zeroize command to zeroize a single Routing Engine on a dual Routing Engine device, the CLI incorrectly displays a message indicating that it will zeroize both Routing Engines.
- **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](#).]

- **Validation of /vm-primary mount during JDM installation or upgrade (Junos Node Slicing External Server Deployment)**—During installation or upgrade of the Juniper Device Manager (JDM) on external servers running Junos Node Slicing Release 25.2 or later, an issue occurred with the validation of the /vm-primary mount that stores the GNF images. When /vm-primary was mounted using logical volumes (LVM), JDM would fail to detect that the underlying storage was an SSD. This issue is now fixed. However, the fix introduces a new dependency on the LVM2 package in the host OS. This package is included by default in standard installations of both RHEL and Ubuntu external servers. However, it is advised that you check if the LVM2 package is already installed on the host before installing or upgrading JDM.
- On the MPC7E-10G line card, when you configure the 10-Gigabit Ethernet ports to operate as 1-Gigabit Ethernet ports, use the speed statement at both the edit interfaces interface name together-options and edit interfaces interface name hierarchy levels.
- **Control Maximum 802.1X Client Connections per Interface**—By default, dot1x interfaces configured in multiple supplicant mode have a client limit of 100 authenticated connections per interface. Any additional connection attempts beyond this limit will be automatically blocked.
- **log-tag functionality:** The log-tag functionality is introduced in [set services service-set]
- **Default route installation for non-default routing instances with iked process (MX480 and MX960)**—You can install default route when the st0 interface is in a non-default routing instance. This

enhancement supports migration from MS-MPC to SPC3 as MX-SPC3 injects these routes through auto route insertion (ARI) for traffic selector routes. The configuration facilitates route installation in specified routing instances.

[See [Traffic Selectors in Route-Based VPNs](#).]

- **In on-device packet capture, the *self-mirror-start* value range now begins at 45 seconds (MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004)**—For on-device packet capture, you set a *self-mirror-start* value that determines the duration of the packet capture. The value range now starts at 45 seconds; it previously started at 1 second.

[See [request forwarding-options port-mirroring instance family self-mirror-start](#).]

- **Static flowspec prefix validation**—Commit validation now rejects static flowspec destination and source prefixes that use a shorter mask than the configured address family permits. The configuration must use a valid prefix length before the commit can succeed. This alignment with static route validation prevents invalid flowspec entries from being installed and ensures consistent routing policy enforcement across IPv4 and IPv6 configurations. [PR1903233](#)
- **Enhanced ping command XML output (Junos OS and Junos OS Evolved)**—The XML output for the CE-IP ping has enhanced details to differentiate between successful and failed ping attempts. The update introduces new tags that track the total number of packets sent and received, while also highlighting host unreachable replies and the overall packet loss percentage.

[See [Supported and Unsupported Features for CE-IP Ping](#).] [PR1906967](#)

- **Pixel-Tracing Command Updates**—The pixel-tracing command suite has been updated to include the following options: flush, disable, and enable. These options allow you to control and manage pixel-tracing processes for routing tasks. Note that the rib-sharding option is removed from the CLI commands to prevent system crashes when sharding configuration is present in the system. Ensure that your configurations do not rely on unmentioned options to avoid potential disruptions in your workflow. [PR1909594](#)
- **request jnu role delete Command Unhidden (Junos OS and Junos OS Evolved)**—The request jnu role delete command, which was previously hidden, is now available for use. This command allows users to delete the role of controllers or satellites directly through the CLI. This enhancement improves administrative control and flexibility in role management within the network infrastructure.

[See [request jnu role](#).] [PR1930291](#)

- **Run CSDS CLI commands as root user (MX240, MX304, MX480, MX960, MX10004, and MX10008)**—Run CSDS CLI operations as the root user to ensure successful execution of maintenance commands for key exchange, authentication, and JDM interactions. Run the following commands as root user: request csds jdm add csds-instance-id csds-instance-id image vsrx-image-with-path request csds authenticate-host csds-instance-id csds-instance-id request csds delete-vsrx csds-instance-id csds-instance-id request csds extract-vsrx-keys csds-instance-id csds-instance-id request

csds jdm add csds-instance-id csds-instance-id image vsrx-image-with-path request csds sync-controller other-controller-ip other-controller-ip If a non-root user runs these commands, you'll see the following warning message - warning: Must be logged in as root to execute command. See [request csds add-vsr, request csds authenticate-host, request csds delete-vsr, request csds extract-vsr-keys, request csds jdm, request csds sync-controller]. [PR1918503](#)

MPLS

- **Label field naming update in show route forwarding-table extensive output (Junos OS and Junos OS Evolved)**—The CLI output for MPLS forwarding entries now uses Label instead of Load balance label. This name change standardizes the field name and improves overall clarity when reviewing MPLS forwarding information. The change affects only the output text and does not modify label allocation, forwarding behavior, or traffic handling. No configuration changes are required. [PR1899543](#)

Network Management and Monitoring

- **Ephemeral database default commit synchronize model changed to synchronous (EX2300, EX2300-MP, EX2300-C, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-48MP, EX4000-24P, EX4000-24T, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, and MX10008, MX10016)**—We've changed the default commit synchronize model for the ephemeral database from the asynchronous model to the synchronous model. With this change, we've deprecated the allow-commit-synchronize-with-gres statement and only the synchronous model supports synchronizing ephemeral data on devices that have graceful Routing Engine switchover (GRES) or nonstop active routing (NSR) enabled.

[See [Understanding Ephemeral Database Commit Synchronize Models](#).]

- **Deprecation of shell option**—The shell option no longer requires a separate configuration and is now the default behavior. Deprecating the shell option enhances efficiency and simplifies management tasks.
- **Shell Command Logging Enhancement**—All shell commands executed on the device and root sessions are now logged by default. This enhancement ensures enhanced security and auditability by capturing all commands entered in any shell environment, preventing bypassing audit logging. [PR1867216](#)

Platform and Infrastructure

- **Commit validation for unique user IDs**—We have added support to validate the user configuration to ensure that each user is assigned a unique UID. A commit fails if duplicate UIDs are detected, ensuring stronger validation and preventing identity conflicts. Previously, a commit was successful even when multiple users shared the same UID, triggering only a warning and logging a syslog message.[PR1890656](#)

Routing Protocols

- **Behavior change for multihop eBGP sessions (MX960)**—When a multihop eBGP session with a directly connected peer goes into IDLE state due to no local interface being present, the system now does a hard reset instead of handling it as a configuration change. This change ensures that the session and associated routes are reset promptly, improving the timeliness and accuracy of routing updates. This adjustment is important for maintaining the stability and efficiency of your network routing, especially in scenarios involving interface deactivation or removal.[PR1869927](#)
- **SRv6 MSD advertisement compliance**—MSD type 43 (SRH Max T.insert) is no longer advertised in IS-IS SRv6 node capability TLVs to match the latest standards draft. This alignment maintains interoperability with peer devices and tools that expect the reduced MSD type list and ensures SRH limits are reported consistently across the network.[PR1879051](#)
- **Auto BGP peer discovery for unnumbered interfaces of iBGP peers**—Junos supports single-hop iBGP sessions for unnumbered peering. Previously, the iBGP implementation created an indirect next-hop for route resolution, even if the iBGP peer was directly connected.[PR1884861](#)
- **BGP route-refresh warning messages added**—New messages have been introduced to the `show bgp diagnostics warnings` command output. These messages provide diagnostic information when the number of BGP route-refresh messages exceed predetermined thresholds within specific time intervals of 5 minutes, 30 minutes, and 1 hour.

[See [show bgp diagnostics warnings](#).][PR1909467](#)

- **BGP UPDATE messages must have the marker field set to all ones**—The marker field in BGP header is a 16-octet field included for compatibility that must be set to all ones. All Junos OS and Junos OS Evolved devices send a notification and tear down the BGP session if the marker field in the header of the received update message is not set to all ones.[PR1955678](#)

Subscriber Access Management

- **Addition of message-authenticator and no-message-authenticator attributes under access radius-server, access profile radius-server, and system radius-server hierarchies**—Set `message-authenticator` if you require the RADIUS server to include the Message Authenticator attribute in replies to Access-Request messages. Set `no-message-authenticator` to not require that attribute.[PR1871147](#)

User Interface and Configuration

- **Updated Annotate Command Behavior**—The annotate command now correctly handles multi-line comments that start with `#`. This change ensures that all lines of a multi-line comment remain comments, preventing unauthorized configuration changes when the configuration is committed, reloaded, or rolled back. These improvements enhance security by mitigating potential privilege escalation attacks from users with limited configuration editing permissions.[PR1868636](#)
- **Enhanced Permission Checks for Rename/Copy Operations**—New permission checks have been introduced for rename and copy operations within the configuration hierarchy. These checks parse the hierarchy being modified to ensure the user has the required permissions for the hierarchy and its sub-hierarchies. If permissions are insufficient, the operation will fail and a "Permission denied" error will be displayed. This enhancement ensures that configuration modifications are performed only by authorized users, improving security and preventing unintended changes to critical system settings.[PR1882303](#)
- **Enhanced file archive exclusion capability (PTX10008)**—The "file archive" operational command now supports ?exclude? option where you can specify a folder or regular expression to exclude from final archive. (Example, exclude traces from /var/log). This enhancement improves the flexibility and efficiency of the archiving process, allowing you to manage log files more effectively. To exclude folders, you can use one of the below formats: user@router> file archive source /var/log destination /var/tmp/abc exclude jam* exclude wtmp* exclude ksyncd exclude aaasd user@router> file archive exclude [license* shmlog ksyncd] source /var/log/ destination /var/tmp/abc[PR1884901](#)
- **Stale ui-state.db data in persistent NETCONF sessions post-mgd restart**—Existing NETCONF sessions might fetch stale data from ui-state.db after mgd -N restart. New sessions correctly map the refreshed database. Scripts must establish new sessions post-restart to access updated values. Functional configuration remains unaffected. [Script failures monitoring "local-host" NETCONF sessions]-Scripts might fail when including "local-host" NETCONF sessions in monitoring operations. Internal sessions are now excluded from tracking. Scripts must filter out "local-host" sessions. No impact to internal application functionality.[PR188557](#)
- **Generate genstate YANG modules on Junos devices**—You can use the show system schema operational command or equivalent RPC to generate the genstate YANG modules in the specified output directory on a device.

[See [show system schema](#).][PR1910388](#)

Known Limitations

IN THIS SECTION

- [General Routing | 91](#)
- [Class of Service \(CoS\) | 91](#)
- [Infrastructure | 92](#)
- [Routing Protocols | 92](#)
- [User Interface and Configuration | 92](#)

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

- On MX10008 and MX10016 platforms with JNP10K-LC480 Line Card installed, the remote end of the Ethernet link of the JNP10K-LC480 Line Card goes down when an ISSU (Unified In-Service Software Upgrade) is performed from Junos OS 24.4R2 (or earlier) to Junos OS 25.2R1 (or higher). This causes traffic disruption. [PR1880150](#)
- 1. Exit the CLI and start the session again. 2. Enter configuration mode and then return to operational mode. 3. Run the 'restart management' command. [PR1890951](#)
- NETCONF Server reply to get operation pointed to /components/component[name=""]/state/will not include any OPTICAL_CHANNEL type component data. [PR1897729](#)

Class of Service (CoS)

- When an AE is configured in HCOS mode and COS traffic-control-profile configuration is also attached to IFLSETs of that AE in a single config commit operation, following error log can be seen on the FPC hosting AE member IFDs: [Error] COS SCHED : CosSchedNode::create - adding child node on parent active token:6549 child node: type:IFLSET nodeIndex:4 parent schedMode:Port parent maxLevels:1 node schedLevel:2 The error log appears because in the log, the "parent maxLevels" value is less than the "node schedLevel" value. This happens because AE is still in a transient state moving to HCOS mode when IFLSET config message is also received simultaneously by the FPC software. However once AE has moved to HCOS mode, the IFLSET config message is replayed by the

software. The software states are set correctly then and no functional issues are seen. This is thus a false alarm in the given scenario. Issue is applicable for all MX Trio AFT based FPCs - MPC10, MPC11, LC9600, LC4800, MX304. If we want to avoid the error log, we can do a 2-step commit where we activate HCOS on AE interface first and commit, then later activate COS config on AE iflset and commit. [PR1905468](#)

Infrastructure

- When upgrading from releases before Junos OS Release 21.2 to Release 21.2 and onward, validation and upgrade might fail. The upgrade requires using the 'no-validate' option to complete successfully. <https://kb.juniper.net/TSB18251> [PR1568757](#)

Routing Protocols

- "advertise-external" feature control-knob behaves at a different scope when a Junos router is acting as pure-PE or RR/ASBR mode for a VPN family. In pure-PE mode, the knob works on per advertised-to-group basis. i.e. an inactive EBGp-route is advertised to a BGP group only if advertise-external is enabled for that group. In RR/ASBR mode, the knob works on global basis, i.e. if it is enabled for any group, an inactive EBGp route in <vrf>.inet.0 is leaked to bgp.l3vpn.0 with the VRF's RD and (assuming unique-RD in use) it is a candidate for advertisement to all BGP groups, export-policy permitting, even if advertise-external is not enabled for those groups. Because there is no path-hiding in bgp.l3vpn.0 table due to unique-RD, advertise-external feature does not further control advertisement of the route to each BGP group. [PR1880531](#)

User Interface and Configuration

- VMX file copy sftp reports error ssh: Could not resolve hostname sftp: Name does not resolve . file copy sftp://10.85.211.4/home/labroot/junk.txt . ssh: Could not resolve hostname sftp: Name does not resolve error: file-fetch failed error: could not fetch local copy of file. [PR1869005](#)

Open Issues

IN THIS SECTION

- [General Routing | 93](#)
- [High Availability \(HA\) and Resiliency | 97](#)
- [Interfaces and Chassis | 97](#)

- Layer 2 Ethernet Services | 97
- MPLS | 98
- Network Management and Monitoring | 98
- Platform and Infrastructure | 98
- Routing Protocols | 98
- User Interface and Configuration | 99

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

- With Next Generation Routing Engine (NG-RE), in some race conditions, the following interrupts messages might be seen on primary RE: kernel: interrupt storm detected on "irq11:"; throttling interrupt source. [PR1386306](#)
- In some NAPT44 and NAT64 scenarios, Duplicate SESSION_CLOSE Syslog will be seen. [PR1614358](#)
- Multiple vulnerabilities have been resolved in MQTT (Message Queuing Telemetry Transport) included with Junos by fixing vulnerabilities found during external security research. Please refer to <https://supportportal.juniper.net/JSA71655> for more information. [PR1651519](#)
- EX4100-24mp,48mp,24p/t,48p/t,F-24p/t,F-48-p/t: In an interop scenario, when using 1G SFP Optic on PIC-2, auto-negotiation should be disabled on the peer. [PR1657766](#)
- 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](#)
- In Netconf private edit configuration session, commit RPC fails when unprotect operation is performed. [PR1751574](#)
- Junos (JET) telemetry that is pre-gNMI telemetry that uses sensors that are of a double data type are converted to a float data type when streamed to a collector. [PR1777319](#)
- JDI-RCT:M/Mx: IS-IS session over MPC11 cards flapped due to "3-Way handshake failed" during ISSU (FRU upgrade stage - reboot phase). [PR1809351](#)
- Zeroize command not working. [PR1857029](#)

- On rebooting DCI (Data Center Interconnect) Gateway device, while the device is coming up, multicast traffic drop is observed in the highly scaled configuration. The traffic drop is observed in the following condition. The rebooted device was the EVPN DF on I-ESI and after coming up, the DF election is triggered and the device is elected as EVPN DF on I-ESI. After coming up and elected as DF, the device builds the multicast routes afresh. This results in traffic drop. [PR1872219](#)
- When VMHost image on both routing engines are installed and rebooted simultaneously, the EEPROM read issue is observed intermittently in RE1. This leads to CB driver not being loaded in RE1 and further leading to LCMD crash. [PR1879559](#)
- Logs from internal ethernet links monitoring script keep repetitively logged to /var/log/messages file if set system syslog file messages user any configuration is used. [PR1886633](#)
- Phase jump of around 300ns is seen upon LAG member switchover from secondary to primary on 400G interface with ZR optics. [PR1893122](#)
- Issue specific to scenario of system/linecard reboot with PTP available but SyncE not available. Otherwise issue is not seen. [PR1897460](#)
- On all Junos MX platforms that have MS-MPC or MS-MIC service cards installed, the use of the CPU throttling can cause the production service sessions to be dropped. [PR1899178](#)
- If the admin tries to configure Proxy-ID and traffic-selector based IPSec tunnels for the same IKE peer(same IKE gateway) and if the proxy-ID(Any-Any IPv4/IPv6) based tunnel gets negotiated first, then the IPSec tunnel corresponding to traffic-selector will not come up as there already exists another tunnel (Proxy-ID) in the system which can also protect the traffic which was intended to be protected by the traffic-selector based IPSec tunnel. Hence, it is recommended that if the admin wants both proxy-id and traffic-selector based tunnel on the system then, the admin should configure unique IKE gateway objects for both of them i.e. unique IKE local-remote gateway pair which means that proxy-ID and traffic-selector should have their individual unique IKE SAs negotiated. [PR1900529](#)
- The EP IFD/IFL output traffic stats as seen via ?show interfaces EP IFD detail? can be different from the EP IFL queue stats as seen via ?show interfaces queue egress EP IFL?. This is because EP IFD and IFL traffic stats are maintained on the EP port's forwarding topology which gets executed prior to the actual queuing. Hence the packets dropped in COS queues are not considered in EP IFD/IFL traffic stats and it shows the queued stats value (not the transmitted stats value) as output stats. Currently there is no support for 'accurate-stats' kind of functionality for EP ports and hence this behavioural difference as compared to non-NGPE WAN interfaces. [PR1901790](#)
- In BGP-CT scenario at ASBR instead of swap operation, we have a pop and push NH programmed which results in pops the transport and service label and then pushes only transport label. Due to this service label is lost and once it reaches Penultimate Hop Router we pops (PHP) the transport label and sends plain IP packet and because service label is lost the DUT is unable to identify the VRF and results in default route reject. [PR1902144](#)

- A vmcore may be observed on the backup when any of the following configurations are performed on the EP IFDs with GRES configured 1. Port Speed change on the EP IFDs When configuring(set)/unconfiguring(delete) port speed on the EP IFDs using the config stanza "chassis port-extender fpc-slot <> pic-slot <> port <> speed <>" 2. Changing the configuration attribute "cascade-port"/"target-mode"/"device" under the config hierarchy "chassis port-extender <> fpc-slot " "chassis port-extender <> fpc-slot cascade-port " "chassis port-extender <> fpc-slot target-mode " "chassis port-extender <> fpc-slot device " Impact: 1. Traffic will be impacted. 2. Backup RE is not available for graceful switchover until state resync/replay completes from the primary. Though vmcore is observed, backup does not reboot as it is system generated live core for debugging. [PR1902701](#)
- On all Junos OS MX platforms that support MPC2E-NG,MPC2E-3D-NG-Q,MPC3E-NG and MPC3E-3D-NG-Q, the Auto-Negotiation (AN) process on certain PHY interfaces of the MIC (MIC-3D-10GE-SFP-E) may intermittently get stuck, preventing link establishment and causing traffic loss. This issue can be triggered by reinserting an SFP-T module, multiple times restarting the mic or by interface driver resets, which lead to inconsistent enable/disable sequences during Auto-Negotiation. [PR1906675](#)
- On MX301 and MX304 in certain configuration scenarios for tunnel and inline services PFE may not be able to allocate streams causing traffic black hole for inline or tunnel services.[PR1907685](#)
- On MX304 with PTP PPM feature enabled, the PTP port state is not moving to Passive after disabling and enabling the interface. [PR1910401](#)
- We have an Issue in NH when we push the configuration, so to work around this we need to activate or deactivate the service-set configured on the interface.[PR1910673](#)
- On Junos and Junos Evolved where Segment Routing Traffic Engineering (SRTE) Tunnel is supported, the log message (RPD_SPRING_TE_ENTROPY_UNSUPPORTED_FOR_ONE_LBL: SPRING-TE Entropy-label is not supported for segment-list with single label for tunnel) is observed even when entropy label feature is not configured. [PR1911821](#)
- As part of the RSI process, the command show vmhost support-info is executed, which comprehensively collects vmhost logs using various cat commands. Some of these commands attempt to access files that do not exist on the MX304, leading to the display of error messages. [PR1913540](#)
- On Junos MX platforms and Junos OS Evolved ACX platforms that support BNG CUPS (Broadband Network Gateway Control and User Plane Separation), L2TP (Layer 2 Tunneling Protocol) subscriber sessions will fail to come up. When this happens, the subscriber session does not bind correctly and no data traffic will pass for the affected subscribers.[PR1916030](#)
- On MX platforms like with AFT based Packet forwarding Engine (MX304/MX301/MX10003) and MX Chassis based systems with AFT based LC there is a limitation of maximum of 1024 ifls on 1G interfaces.[PR1921080](#)
- While verifying export EFP group, sensor components verification failed[PR1922040](#)

- On MX301 rpd may core when logical-systems configuration is used to create EVPN-VPWS [PR1922759](#)
- When Macsec is enabled on 1G interface, Auto-neg needs to be disabled for MX301. Config to disable auto-negotiation: set interfaces <1g-intf-name> gigether-options no-auto-negotiation [PR1922797](#)
- Starting 25.4R1 , Flood traffic in EVPN MPLS deployment on the NodeSlice platform, might experience drop for certain flows [PR1922876](#)
- there is the change tin log message for the CB 0_TEMP_0 sensor while testing RCB on MX301 [PR1924853](#)
- traffic drop seen for VPLS Unicast traffic over IRB over LSI with enhanced-convergence knob configured. [PR1925479](#)
- l2ald Core @l2ald_gbptag_change_mac_in_db, @gbptagdb_update_internal_taginfo is seen in Suite teardown [PR1925496](#)
- On MX10004 and MX10008 devices, display-only issue in Junos CLI show chassis environment : current/power for some of the POLs are shown as 0, observed in 25.2R1-S1. Fixed in later releases. [PR1916094](#)
- Humidity Sensor CLI command is not applicable on MX10004 and MX10008 devices. The command has been suppressed in later releases. [PR1909435](#)
- Support for Virtium SSD firmware upgrade on MX10004 and MX10008 devices not available in Junos OS 25.2R1-S1 Release. Fixed in later releases. [PR1907227](#)
- During ISSU, some traffic drop may occur in HW sync phase for few L3VPN flows. The traffic recovers post HW sync phase. This issue is not specific to 25.4R1. [PR1925599](#)
- HTTPv6 redirect functionality does not work on the MX301 platform. [PR1925736](#)
- On MX301 after Small Form-factor Pluggable transceiver (SFP) Online Insertion and Removal (OIR) on a Media Access Control Security (MACsec) enabled interface, the link and MACsec session remain UP, but MACsec encryption and decryption stop functioning. This results in traffic loss and causes service impact. [PR1932925](#)

High Availability (HA) and Resiliency

- Graceful Routing Engine Switchover (GRES) not supporting the configuration of a private route, such as fxp0, when imported into a non-default instance or logical system. Please see KB <https://kb.juniper.net/InfoCenter/index?page=content&id=KB26616> resolution rib policy is required to apply as a work-around. [PR1782934](#)

Interfaces and Chassis

- When a GE interface that is part of the interface-set is added as a member of an AE interface, and if the AE interface units are added to the same interface-set in which the GE interface was present in the previous commit, then an error message can be seen in the logs mentioning about a failure in deleting the IFL from interface-set. This is happening due to a race condition between creation of the AE interface and the updating the AE IFLs to the interface-set. The exact sequence of operations is as below (the interface mentioned below can be XE/GE/FE/ethernet interfaces that can be added to an AE bundle.)
 commit 1 ----- - Create IFLs on a GE interface. - Add the GE IFLs to an Interface set, say iflset1
 commit 2 ----- - Delete the GE interface config - Create AE interface with units. - Add the GE interface as member of AE interface - Add the AE IFLs to interface set iflset1 (same interface set in commit 1) - Delete the GE interface from interface set iflset1
 When the above config changes are committed, error message similar to the below could be seen in syslog when the issue is hit
 [Error] IF:IfI not deleted, not present in IfISet, ifISetIndex:2 ifIIndex:355
 This has no impact on functionality or stability of the box. The workaround for the same could be to split the commit 2 as -
 Delete the GE interface config - Delete the GE interface from interface set iflset1 - commit - Create AE interface with units. - Add the GE interface as member of AE interface - Add the AE IFLs to interface set iflset1 (same interface set in commit 1) - commit
 Also please check the other limitation related to CoS and Aggregated Interface mentioned in the link below <https://www.juniper.net/documentation/us/en/software/junos/cos/topics/concept/schedulers-cos-ae-sdh-limits-cos-config-guide.html> [PR1905458](#)

Layer 2 Ethernet Services

- On all Junos devices supporting subscriber services, in case of dual stack DHCP (Dynamic Host Configuration Protocol) subscribers with IA_NA (Identity Association for Non-temporary Address) and IA_PD (Identity Association for Prefix Delegation) bindings with lease times (For the assignment of IPv6 address to a client device), when a client initiates separate renew exchanges for the IA_NA and IA_PD, and once client and DHCP server are in sync with these timers, there can be a race condition at Junos device which is DHCPv6 relay, has not refreshed lease timer and can go out of sync. This can result in deleting IA_NA/IA_PD binding and route to get deleted for that subscriber only. This causes one of the leg for IA_PD or IA_NA to go down for that subscriber, which can result in traffic impact for that leg. [PR1911001](#)

MPLS

- On MX301, rpd core might be seen after deactivate/activate protocol IS-IS trigger [PR1901973](#)

Network Management and Monitoring

- On MX301 snmd might core in certain scenarios. [PR1913131](#)

Platform and Infrastructure

- With HCOS 2 level hierarchy configuration in logical tunnel interface, the forwarding pipeline for LT interfaces was not having the right queue information leading to the packets getting stuck in memory which leads to an increase in usemeter numbers causing the FPC get into major error. [PR1767970](#)
- An Authentication Bypass by Spoofing vulnerability in the RADIUS protocol of Juniper Networks Junos OS and Junos OS Evolved platforms allows an on-path attacker between a RADIUS server and a RADIUS client to bypass authentication when RADIUS authentication is in use. Please refer to <https://supportportal.juniper.net/JSA88210> for more information. [PR1850776](#)
- There are limitations associated with the implementation features of the generator in the Broadcom chip. Multiple Y.1564/RFC2544 generators should not be used on the same interface. Juniper engineering recommend running tests one by one or distributing tests across different VLANs. This PR will commit code to ensure an error is presented when trying to run multiple tests. [PR1908499](#)
- On MX301, IGMP might not get enabled by default when PIM is enabled. Use `set protocols igmp interface` to explicitly enable IGMP. [PR1918590](#)
- On MX301, the `fetch` command verbose output has changed causing scripts to fail. As a workaround, you must change the expect script matching string. [PR1924159](#)

Routing Protocols

- This is system limitation due to high system load and aggressive IS-IS hello timer. Workaround: hello Timer need to increase in order adj not to flap. [PR1314650](#)
- On all Junos platforms, after interface configuration rollback, sessions stay in Idle state when multiple BGP(Border Gateway Protocol) sessions exist. [PR1880630](#)
- In large scale routing instance configuration, please provide interval of 20-30 mins for system to stabilize. [PR1883895](#)
- Resolver optimizations for BGP labeled families will work if the "per-prefix-label" allocation mode is used. When using per-nexthop label allocation, the routes resolution will not be optimized in releases where this PR is not committed. Workaround has been provided. [PR1900514](#)

- On MX301, IGMP might not get enabled by default when PIM is enabled. Use `set protocols igmp interface` to explicitly enable IGMP. [PR1918590](#)

User Interface and Configuration

- Getting validation error for get-interface-information rpc execution through ODL controller. [PR1899597](#)

Resolved Issues

IN THIS SECTION

- General Routing | [100](#)
- EVPN | [105](#)
- Forwarding and Sampling | [105](#)
- J-Web | [105](#)
- Juniper Extension Toolkit (JET) | [105](#)
- Layer 2 Features | [106](#)
- MPLS | [106](#)
- Network Management and Monitoring | [106](#)
- Platform and Infrastructure | [106](#)
- Routing Policy and Firewall Filters | [107](#)
- Routing Protocols | [107](#)
- Services Applications | [108](#)
- Subscriber Access Management | [108](#)
- User Interface and Configuration | [108](#)
- VPNs | [108](#)

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.

General Routing

- Adjust gNMI SetRequest Update handling for list keys + value payload. [PR1691474](#)
- IS-IS Level 2 disabled after upgrade due to OpenConfig YANG model change. [PR1798178](#)
- The "user.notice logrotate: ALERT exited abnormally with [1]" messages can be seen on a system with an MPC10E. [PR1833493](#)
- CTRL FPGA Version Check for MX10K-LC480. [PR1846388](#)
- Loss of synchronization are observed when SyncE/PTP configurations are changed. [PR1848235](#)
- After GRES ptp-spll-lock-state is stuck in ACQUIRING instead of phase aligned. timingd re core is also observed which is not seen in rerun. [PR1855379](#)
- Deprecate a few sensor paths that are not OpenConfig compliant. [PR1858680](#)
- AE Re-anchoring Impacting CFM on untagged AE interfaces. [PR1864491](#)
- On MX Platforms, FRU Insertion/Removal Traps for Fan are not generated. [PR1866649](#)
- Asymmetry configuration on slave ports causes small spikes. [PR1867423](#)
- L2ald memory usage increased up to 100MB. [PR1867811](#)
- PPE errors will be seen during ISSU. [PR1868224](#)
- Subscribers failed to establish DS-Lite softwires due to stale softwire entries. [PR1869450](#)
- CUPS-L2-Regression: dbng_single_up_igmp_mld_mts_part1_mts.robot fails in teardown with error no response from radius server. [PR1871403](#)
- Restarting the evo-aftmand-bt or evo-cda-bt process disrupts PHY synchronization within the timingd service resulting in timing errors. [PR1878029](#)
- The out-of-order delete messages are seen after core facing interface with SCU configuration flaps. [PR1878057](#)
- Link flapping occurs on stable subinterfaces when using QDD400GZR optics if any one subinterface in the channelized group is down. [PR1878198](#)
- GNFs in Junos Node Slicing will stop responding when both REs are rebooted at the same time and primary RE fails to recover. [PR1880218](#)
- NAT Pool Installation failure due to Service-Set name length mismatch. [PR1881192](#)
- On MPC11E linecards, Periodic error messages are logged as **Temp sensor DDR4 A failed**. [PR1881499](#)

- MPC10E: Huge 1PPS error is seen for 15 to 20 minutes with SCB3E post DUT in phase aligned state. [PR1881618](#)
- Firewall policy configured to match IP payloads fail matching on MPLS packets. [PR1882315](#)
- The em0 mgmt port is unreachable after RE switchover. [PR1882329](#)
- BFD fail to establish over an IPsec tunnel on Juniper MX Series with the SPC3. [PR1882490](#)
- EVPN-MPLS BUM Traffic Disruption Due to Incorrect QinQ STag Insertion. [PR1882561](#)
- MPC 10/11/12E, LC9600 and LC4800 Line cards and MX304, interface statistics stop after interrupt. [PR1882845](#)
- The debug-collection fails due to insufficient space. [PR1883317](#)
- Packet drops are observed when MACsec with bounded-delay is configured due to key rollover. [PR1883473](#)
- FPC crash is seen on certain Junos platforms when firewall filter is configured for subscribers. [PR1883530](#)
- AE interface going down on specific MX platform with exclude-protocol being re-configured under MACsec. [PR1885185](#)
- Application rpd-agent might restart with a coredump after interface related event changes. [PR1885455](#)
- TTE miscalculates bandwidth on AE interfaces due to zero negotiated bandwidth. [PR1885563](#)
- IGMP or MLD packets associated with CCC services will be dropped instead of being forwarded. [PR1885670](#)
- MX304 LNS: FPC restart and aft-trio core after LMIC OIR when SI pool spans both MICs. [PR1885754](#)
- Few telemetry paths are not exported after router reboot. [PR1886043](#)
- The debug_collector: support multiple modes data collection. [PR1886371](#)
- Alarms for high usage in /var partition are not generated. [PR1886757](#)
- Incorrect tunnel and SA count display due to simultaneous IKEv1 quick mode negotiation. [PR1886887](#)
- The snmp mib walk DomCurrentLaneAlarms does not show proper lanes Values in the latest builds. [PR1886889](#)

- Interfaces either fail to come up or flap or a delay is observed on MX10003 platforms when the interface is reset or the devices is restarted. [PR1886937](#)
- Optics fail to come up post reboot. [PR1887528](#)
- Packet loss observed with SFP-T modules on MX10K LC480 line cards due to IPG misalignment. [PR1887864](#)
- Incorrect uSID handling in SRv6-TE will impact the traffic path. [PR1887866](#)
- Change in the behaviour of configured lifeseize kilobytes of ipsec proposal. [PR1888205](#)
- Temporary file /tmp/veriexec_test is not cleaning up on all platforms after "request system malware-scan integrity-check". [PR1889037](#)
- Traffic drop is observed in an EVPN multihoming as the MAC route points to the ESI interface when the CE (ESI) IFD flaps. [PR1889335](#)
- DHCP clients do not come up when VRF leak and "dhcp-relay" with "no-snoop" are configured under a routing-instance. [PR1889637](#)
- Data still seems to be streaming somewhere when the DialOut (Established) connection on the port is already closed. [PR1889924](#)
- BNG CUPS Controller will not deploy on a Multi-Geo RHOCN Multicenter. [PR1890548](#)
- Mismatch between configured and reported wavelength for some optics. [PR1890558](#)
- Enable high-power optics on all ports, limited to a maximum of 32 high-power modules. [PR1890645](#)
- SFB is stuck in CMTY_SFB_STATE_OFFLINE_ACK_WAIT state and does not come offline or online. [PR1891047](#)
- A GRE tunnel configured with a tunnel key drops MPLS-encapsulated traffic. [PR1891110](#)
- The hardware-assisted-keepalives CFM transmit packets are queued into Q4 instead of Q3/ Network-control. [PR1891363](#)
- On MX304 platform ungraceful removal of standby RE results in communication loss between REs. [PR1891577](#)
- Clear security log reports with time interval support. [PR1892154](#)
- Incorrect shared memory unit interpretation causes memory leak log errors. [PR1892348](#)
- Memory leak and L2ALD process crash caused by SNMP polling of Q-BRIDGE-MIB on devices configured with RTG interfaces. [PR1892944](#)
- Adding a new key to authentication-key-chain causes kernel crash. [PR1895827](#)

- MX[10008, 304] Marvell Phy: MACSEC:TRD: With should-secure enabled and IFL interface configured with IFD level MACsec, ping is not working. [PR1896363](#)
- SFP-T port will not come up after system restart. [PR1896458](#)
- The transportd daemon which runs on Routing Engine cards generates core dump repeatedly after installing line cards or PICs are swapped for a different type. [PR1896486](#)
- Incorrect ARP responses observed in BNG CUPS software versions. [PR1897105](#)
- Traffic loss occurs with multiple VRRP instances using the same VRRP group ID and one of VRRP instance transitions backup. [PR1897117](#)
- ARP resolution and device discovery failure is observed due to unexpected VLAN tags on ARP replies. [PR1897336](#)
- Memory allocation failure in all the FPCs inside the NH partition. [PR1897464](#)
- PFEs fail to come online after a rapid restart or power cycle. [PR1898307](#)
- The craft-control process is unable to start on MX10004 and MX10008 platforms. [PR1898722](#)
- The "gnsi_ssl_profiles" file is not getting synchronised between the routing engines. [PR1898768](#)
- The vMX line card is crashing due to internal debug logging issues. [PR1898885](#)
- MAC learning failure when moving the AE interface from one VLAN to another VLAN in a single commit. [PR1899530](#)
- Unexpected packet retransmissions and traffic drops observed on SFP-T (1G copper) interfaces on LC480 line cards. [PR1899711](#)
- In Fusion Provider Edge deployments, aggregator does not establish LLDP neighborship with satellite interfaces. [PR1900111](#)
- SPC3 continues to crash after attempting to configure NETFLOW inline-service flow-table-size and mpls-flow-table-size. [PR1900449](#)
- Service-Set Configuration Bug Leading to Kernel Panic on Junos MX. [PR1901021](#)
- Anomalies observed on master RE during switchover. [PR1901341](#)
- FPC unresponsive due to memory exhaustion on PTX platforms. [PR1901728](#)
- [Clocking Solution:MX480]:SyncE stuck in Holdover after PTP protocol alone deactivated from the G.8275.1 profile configuration. [PR1902478](#)
- Intermittent kernel panic results in device reboot or fxpc crash. [PR1902609](#)

- Deactivating or deleting auto-configure configuration statement in an interface is not working properly. [PR1902855](#)
- Vmcore triggers while configuring new members into an existing AMS interface on all MX platforms with specific linecard. [PR1903211](#)
- vJunos-switch and vJunos-router can now be deployed on AMD servers. [PR1903528](#)
- Stale entry in MAC-IP table affects ARP resolution. [PR1905196](#)
- Redundant prefix information is included within one router-advertisement packet when configuring accept-data in VRRPv6. [PR1905553](#)
- Correct values of ACCTG-Terminate-cause are not getting populated in ACCTG-STOP message. [PR1905612](#)
- PPE thread timeout traps during EVPN BUM forwarding due to Invalid weighted-unilist selector. [PR1905718](#)
- FPC crash triggered when a line card reboots with a large number of static subscribers. [PR1905768](#)
- While collecting RSI, takes long time to produce output on MX platform. [PR1906557](#)
- GNF will be in offline state and its mastership state is set to backup on both routing engines when RE is forced to boot from SSD2 or normal RE switchover on MX platforms. [PR1906960](#)
- CLI based Virtium SSD firmware upgrade is missing. [PR1907227](#)
- All Marvell(MX, ACX): LLDP protocol goes down when 'exclude protocol LLDP' is deleted from MACsec policy attached to IFD. [PR1908196](#)
- RIB and the FIB inconsistency results in traffic loss in IPsec scenario with st0 interface configured. [PR1908681](#)
- On all MX platforms chassisid gets stuck and becomes unresponsive it resumes only after restarting both control units. [PR1908719](#)
- Suppress humidity sensor CLI command and corresponding code in MX301. [PR1909435](#)
- Selection of VGA-IP as source IP for RE-ARP packet causes incorrect ARP updation of VGA-IP getting bound with IRB-MAC at end-hosts. [PR1909786](#)
- SPC3 crashes when three-color policer is attached to firewall filter. [PR1910534](#)
- Show command execution failure for show system macsec license on MX platforms. [PR1911538](#)
- The nsd process crash will be seen on MX platforms when configuration change is committed using ephemeral database. [PR1912459](#)

- EVPN routes are stuck in the KRT queue. [PR1913519](#)
- GRE-over-GRE tunnel is down due to keepalive packets dropped. [PR1913870](#)
- On-change telemetry events dropped when l2ald telemetry queue memory limit is not configured. [PR1914423](#)
- Total count will be incremented in show dynamic-tunnel summary for tunnels. [PR1917815](#)
- XML validation failure is observed while verifying. show lldp neighbors and show lldp neighbors interface. [PR1914374](#)

EVPN

- The data plane will be out of sync when migrating to EVPN A/A stitching with Vanilla VXLAN (PIM Multicast). [PR1848993](#)
- The rpd EVPN module sends a redundant license message to the license infrastructure. [PR1889092](#)
- EVPN routes take longer to install into the FIB after being learned through BGP. [PR1893671](#)
- Inconsistency is observed between the ARP table learned on PE devices in EVPN-MPLS or EVPN-VXLAN Multihoming scenario. [PR1894803](#)

Forwarding and Sampling

- Filter name mapping error in routing instances with attached firewall filter. [PR1889028](#)
- Memory leak in FPC during login or logout. [PR1890097](#)
- Intermittent traffic loss after pfe reset due to FLT filters. [PR1903047](#)
- [MX10008] cmd='ls -i /var/etc/filters/filter-define.conf' is logged every 1 second instead of every 30 seconds. [PR1903874](#)

J-Web

- Junos OS: J-Web: Multiple vulnerabilities resolved in PHP software (CVE-2023-0567, CVE-2023-0662, CVE-2023-3823, CVE-2023-3824, CVE-2023-0568). [PR1725808](#)

Juniper Extension Toolkit (JET)

- The master-eventd will fail after multiple RE switchover. [PR1872284](#)

Layer 2 Features

- In VPLS scenario due to ungraceful switchover rpd process crash is observed. [PR1882938](#)
- VPLS PNH configured on GRES enabled system results in traffic loss post the RE switchover. [PR1903609](#)

MPLS

- On rare circumstances the rpd process crash is seen in RSVP scenario. [PR1871664](#)
- MPLS ping/trace not working for direct peers via routing-instance over MPLS protocols. [PR1889546](#)
- Record Route Object displayed in `show mpls lsp` output is truncated if number of hops is sixteen or more. [PR1893822](#)
- More bandwidth admitted onto a TE link when Label Switched Paths (LSPs) undergoing make-before-break re-route over the same link carrying the bypass LSP during local repair. [PR1896022](#)
- Frequent link-protection flaps are observed for container LSP's with no change in member LSP. [PR1908506](#)

Network Management and Monitoring

- Bug in `ddl_expand_child` function adds an extra "-n" flag after comparing the command passed with `(/bin/sh)`. [PR1848295](#)
- Syslog forwarding intermittently stops post DUT reboot on virtual devices. [PR1853209](#)
- JDI-REG:VIRTUAL:EVO: After management-instance is configured, Observing command timeout in device while establishing netconf sessions over outbound https. [PR1865547](#)
- Configuration under `set system trace application <app-name>` gets lost during rollback with certain sequence of steps. [PR1869479](#)
- GNMI get native configuration garbage reply when there is no configuration related to openconfig. [PR1879816](#)
- The `snmpwalk` for OID `ifJnxInputErrors` is not working on Junos Evolved platforms. [PR1890712](#)

Platform and Infrastructure

- DDOS/SCFD culprit-flows display wrong PPS on the detected subscriber demux0 interfaces. [PR1860439](#)
- Traffic flow display not accurate when SCFD is enabled. [PR1897237](#)

Routing Policy and Firewall Filters

- Deleting duplicate entries in a route-filter-list or source-filter-list can disrupt policy evaluation, leading to route withdrawal and traffic disruption for services using the filter-list. [PR1887006](#)
- SRX and MX-SPC3: While Downgrading from 25.4/24.4R2/25.2R2/25.3R1 image with address book IPV4 range config, image installation validation is failing. [PR1899519](#)

Routing Protocols

- Scaled BGP sessions stuck in idle after interface rollback. [PR1880630](#)
- The rpd process crashes after BGP configuration commits involving group-split-size and RIB-sharding. [PR1887911](#)
- BGP Prefix-SID Label collision causing RPD crash. [PR1889749](#)
- Primary path temporarily disappear until the MLA timer expires for IPv6 SRv6 routes. [PR1892136](#)
- The rpd process crashes in an Inter-AS Option-AB L3VPN with BGP multipath list-nexthop enabled [PR1898734](#)
- On platforms supporting BGP RIB sharding the rpd process crash is observed on both REs. [PR1903829](#)
- IS-IS FA configuration in IPv4 or IPv6 Multicast topology might cause inconsistency in routing table leading to routing loop. [PR1906578](#)
- Routes are hidden when accept-own feature is enabled with rib-sharding. [PR1907391](#)
- The rpd process crashes when generate route is configured in a VRF having VXLAN EVPN routes. [PR1907558](#)
- Bgp prefixes get stuck in output queue forever, with bgp delay-route-advertisements and route-ack-converge feature enabled. [PR1909599](#)
- Junos and Junos OS Evolved platforms experience high CPU after FPC reboot causing unpredictable issues with protocols (OSPF/IS-IS/BGP, etc.) managed by PPMD. [PR1909719](#)
- Memory leak is seen during VRF add/delete when vrf-table-label is present. [PR1910691](#)
- BFD sessions will not come up on Junos OS and Junos OS Evolved platforms due to keychain names overlapping. [PR1912250](#)
- BGP task replication will be stuck in 'InProgress' state following an RE switchover when two single hop EBGP sessions are configured on two different interfaces using the IP addresses from the same subnet. [PR1914814](#)

Services Applications

- PAA installation failure on reboot due to "Not found default vrf" error. [PR1886928](#)

Subscriber Access Management

- The 'bbe-smgd' process crashes in BNG with APM environment when loopback interface is moved/ removed from a routing-instance through CLI. [PR1862071](#)

User Interface and Configuration

- The dcd process crashes when deactivating only 'swap' under ' interfaces <> unit <> output-vlan-map' with no other attributes present. [PR1872820](#)
- The mgd process crash is seen on all Junos and Junos Evolved platforms when FQDN is configured along with ephemeral database. [PR1878430](#)
- The xnm-ssl feature fails without loopback interface configuration on Junos Evolved platforms. [PR1882996](#)
- Slow configuration commit observed on devices with a single Routing Engine (RE). [PR1884781](#)
- Device enters Amnesiac mode when converting Routing Engine from non-USF mode to USF mode. [PR1889570](#)

VPNs

- In Dual PIM MVPN setup the rpd process crashes in a rare case scenario. [PR1883458](#)
- MVPN Source PE might incorrectly send mcast traffic on SPT while actual receiver is still on RPTree. [PR1888630](#)
- The rpd process crash is observed with MVPN and RIB sharding enabled. [PR1902405](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 112](#)

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 25.4R1



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

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

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

For more information about the installation process, see [Installation and Upgrade Guide](#) and [Upgrading Junos OS with Upgraded FreeBSD](#).

Procedure to Upgrade to Junos OS

To download and install Junos OS:

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

8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new jinstall package on the routing platform.



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

All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-
x86-32-25.4R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-
x86-64-25.4R1.9-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-
x86-32-25.4R1.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-
x86-64-25.4R1.9-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 21.1R1, 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](#).

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 25.4R1

To downgrade from Release 25.4R1 to another supported release, follow the procedure for upgrading, but replace the 25.4R1 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](#).

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 9: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
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 NFX Series

IN THIS SECTION

- [What's New | 114](#)
- [What's Changed | 114](#)
- [Known Limitations | 115](#)
- [Open Issues | 115](#)
- [Resolved Issues | 116](#)
- [Migration, Upgrade, and Downgrade Instructions | 117](#)

What's New

IN THIS SECTION

- [Application Layer Gateways \(ALGs\) | 114](#)

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

Application Layer Gateways (ALGs)

- **Support for client identifier in forwarded DNS queries using experimental EDNS(0) option (NFX150, NFX250, NFX350, cSRX, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—You can enhance DNS query responses by utilizing the *Client ID in Forwarded DNS Queries* feature on SRX Series devices. This functionality supports precise client identification, essential for services such as parental control. You can configure SRX Series Firewalls to include client identifiers such as MAC addresses, IPv4, or IPv6 addresses in DNS queries. This feature, intended for controlled environments, generates targeted DNS responses based on the originating device's identity, improving service accuracy and network efficiency.

[See [DNS ALG](#).]

What's Changed

IN THIS SECTION

- [Dynamic Host Configuration Protocol | 115](#)
- [Platform and Infrastructure | 115](#)
- [Routing Protocols | 115](#)

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

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

Platform and Infrastructure

- When you configure the security policies on the NFX devices using the CLI, [edit security policies from-zone trust to-zone untrust policy *policy-name* match application], the default application, `junos-vxlan` now appears in the application list.

Routing Protocols

- **BGP UPDATE messages must have the marker field set to all ones**—The marker field in BGP header is a 16-octet field included for compatibility that must be set to all ones. All Junos OS and Junos OS Evolved devices send a notification and tear down the BGP session if the marker field in the header of the received update message is not set to all ones.[PR1955678](#)

Known Limitations

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

Open Issues

IN THIS SECTION

- [General Routing | 116](#)
- [VNFs | 116](#)

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 NFXplatforms 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. [PR1753117](#)
- When there's a change in the VLAN ID from Layer 2 family ethernet switching to Layer 3 family inet, we notice a brief traffic drop during the initial seconds. This behavior suggests that ARP and MAC learning only commenced after the 17th packet, leading to initial packet loss following the VLAN change. Once the ARP and MAC tables were fully populated, packet forwarding resumed normally, and no further loss was observed. [PR1867867](#)

VNFs

- On the NFX350 and NFX250 devices, VNF related SNMP traps are not generated when the client IP is configured. [PR1868397](#).

Resolved Issues

IN THIS SECTION

- [General Routing | 116](#)

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.

General Routing

- On NFX350 platforms, when connected to specific peer devices via SFP 1 Gigabit Ethernet (GE) with auto-negotiation enabled, there may be a status mismatch during the initial handshake. Consequently, the port status on the peer device may display as up/down, which impacts traffic flow.

This issue does not arise when the peer device is either an NFX250 or an SRX-Series platform.
[PR1858495](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 117
- [Basic Procedure for Upgrading to Release 25.4](#) | 118

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 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 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](#).

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

Basic Procedure for Upgrading to Release 25.4

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 25.4R1:

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 | 120](#)
- [What's Changed | 127](#)
- [Known Limitations | 133](#)
- [Open Issues | 133](#)
- [Resolved Issues | 134](#)
- [Migration, Upgrade, and Downgrade Instructions | 136](#)

What's New

IN THIS SECTION

- [Class of Service | 120](#)
- [EVPN | 120](#)
- [Flow-Based and Packet-Based Processing | 122](#)
- [Junos OS API and Scripting | 123](#)
- [Junos Telemetry | 123](#)
- [MAC Learning | 124](#)
- [Multicast | 124](#)
- [Network Address Translation \(NAT\) | 124](#)
- [Post-Quantum Cryptography \(PQC\) | 125](#)
- [Routing Policy and Firewall Filters | 125](#)
- [Routing Protocols | 125](#)
- [Additional Features | 127](#)

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

Class of Service

- **ECN support on MPLS networks (QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**— With this feature, you can configure ECN pairs (A and B) in the EXP header, where A denotes that the packet is ECN capable and B indicates that congestion is experienced. Defining these pairs in MPLS networks enables marking of ECN-capable traffic instead of dropping packets.

To configure the ECN pairs, use the `set class-of-service mpls-ecn-map ecn-capable-exp A congestion-experienced-exp B` command. You must maintain consistent ECN mappings across network nodes for seamless operation.

[See [ECN Support on MPLS Networks](#).]

EVPN

- **Unified access policy (EX4100, EX4400, EX4650, and QFX5120)**—Unified access policy extends group-based policy (GBP) support to Mist APs, including to parts of the wired and wireless access

network outside of the EVPN-VXLAN infrastructure. GBP tags are learned through proprietary control plane messages from Mist APs and across access switches, allowing both wired and wireless clients to participate in GBP microsegmentation.

[See [Microsegmentation Using Group-Based Policies](#).]

- **GBP support for DHCP, ARP, and neighbor discovery packets when snooping and inspection are enabled (EX4100, EX4400, EX4650, and QFX5120)**—DHCP snooping, dynamic ARP inspection, and dynamic IPv6 neighbor discovery inspection now include GBP support for DHCP, Address Resolution Protocol (ARP), and neighbor discovery packets, respectively. Previously, when snooping and inspection were enabled, GBP processing of the snooped and inspected packets did not take place.

[See [Microsegmentation Using Group-Based Policies](#).]

- **GBP on an IPv6 underlay (EX4100, EX4400, EX4650, and QFX5120)**—Group-based policy (GBP) is now supported on top of an IPv6 underlay network. With an IPv6 underlay, you can take advantage of the expanded addressing capabilities and efficient packet processing that the IPv6 protocol offers.

[See [Microsegmentation Using Group-Based Policies](#) and [EVPN-VXLAN with an IPv6 Underlay](#).]

- **EVPN maintenance mode CLI for multihomed ERB leaf nodes (EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200, and QFX5210)**—You can streamline the upgrade process for EVPN-VXLAN leaf devices by utilizing the *maintenance mode* CLI. This feature enables you to isolate multihomed nodes and manage the upgrade with minimal traffic loss. Use the configuration command `set protocols evpn maintenance-mode erb-leaf action-type choice` to enable maintenance mode, and verify the status with `show evpn maintenance-mode status`. Ensure prechecks are validated to prevent disruptions, and manage the process efficiently with provided commands for deletion and validation.

[See [EVPN Maintenance Mode for Multihomed Leaf Isolation](#).]

- **EVPN multihoming and multitenancy support over colored IP fabric with BGP DPF (EX4100-24MP, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—You can leverage EVPN-VXLAN over colored IP fabric using BGP deterministic path forwarding (DPF) to support multihoming and multitenancy configurations for AI/ML applications. This functionality facilitates EVPN for Layer 3 networks with EVPN Type 5, enhancing network segmentation and resource allocation. By using a colored logical fabric, you can achieve flexible routing as uncolored routes integrate seamlessly with all color-coded sessions, optimizing network efficiency and adaptability.

[See [BGP Deterministic Path Forwarding in a CLOS Network](#).]

- **CRB multicast with IGMP snooping (QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—You can enable centrally routed bridging (CRB) multicast in EVPN-VXLAN fabrics to route between VLANs at spine IRB interfaces and constrain traffic with IGMP snooping. On CRB spine devices, use the `set routing-instances instance-name multicast-snooping-options conserve-mcast-routes-`

in-pfe command to conserve forwarding resources by not installing Layer 2 routes in the Packet Forwarding Engine. Do not configure this option on leaf devices.

[See [Multicast Support in EVPN-VXLAN Overlay Networks, Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment](#) and [multicast-snooping-options](#).]

- **EVPN link bandwidth extended community (QFX5120-32C, QFX5120-48T, and QFX5120-48Y)**—You can load-balance traffic for Type 5 routes based on Layer 3 link bandwidth from multihomed devices using the EVPN link bandwidth extended community. Use this community to provide weighted equal-cost multipath (WECCMP) for unequal load balancing. You enable the bandwidth advertisement with a routing policy. To configure, use the set policy-options policy-statement *policy-name* then aggregate-bandwidth and set routing-instances *instance-name* protocols evpn ip-prefix-routes export *policy-name* commands. Associate bandwidth with OSPF or IS-IS using spf-options multipath weighted one-hop and with BGP using the multipath and link-bandwidth auto-sense configurations. If any provider edge (PE) device lacks bandwidth or a units mismatch occurs, forwarding reverts to ECMP. Monitor status with show evpn ip-prefix-database extensive.

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

- **Host route management in scaled, stretched VXLAN campus networks (EX4100-48MP, EX4100-24MP, EX4100-24T, EX4300-MP, 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)**—

We provide new routing policy options, sample routing policies, and new statements to optimize how remote host routes are managed across the access, distribution, and core layers in large-scaled VXLAN campus networks stretched across sites using DCI. With this solution, you can configure the network to install remote host routes in the core layer but not in the distribution and access layers. The distribution and access layer devices only install EVPN Type 5 subnet routes for remote host destinations. The configuration includes policies to ensure the EVPN Type 5 subnet routes are the preferred routes on the distribution and access layer devices. This design reduces the route table burden on access and distribution devices as the network scales.

[See [EVPN Host Route Management for Large-Scale Stretched Campus Networks](#).]

Flow-Based and Packet-Based Processing

- **Unknown unicast drop configuration for VLAN interfaces (QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—You can enhance network performance and prevent traffic storms by configuring your switch to drop unknown unicast packets. This action prevents the flooding of unicast packets with unknown destination MAC addresses across VLAN interfaces.

By default, this feature is disabled. When you enable this feature, the switch learns and adds the source MAC address to the MAC address table. The switch drops packets with unlearned destination MAC addresses. This approach ensures efficient network resource usage and optimal network performance.

To enable the discarding of packets, use the `set switch-options unknown-unicast-forwarding vlan vlan-name drop` command.

To disable the discarding of unknown unicast packets (if you had earlier enabled the feature), use the `delete switch-options unknown-unicast-forwarding vlan vlan-name drop` command.

[See [Understanding and Preventing Unknown Unicast Forwarding](#).]

Junos OS API and Scripting

- Support for libslax 3.1.6 and SLAX version 1.3 (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We've upgraded the libslax library to version 3.1.6, which corresponds to SLAX version 1.3. The upgraded library includes:
 - Slax processor enhancements including a new mode, additional options, and simplified argument parsing
 - New libslax extension library functions
 - Improved SLAX syntax options
 - New SLAX functions and enhancements to existing functions and statements
 - Hexadecimal numbers support in SLAX scripts

This update aligns SLAX processing with contemporary scripting conventions, ensuring efficient, readable scripting while maintaining backward compatibility.

[See [libslax Distribution Overview](#).]

Junos Telemetry

- Enhance VLAN/MAC statistics support for streaming telemetry (EX4300-MP, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48XP, EX4400-48P, EX4400-48T, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214,

QFX5120-48Y, QFX5120-48Y-VC) You can enhance the granularity of streaming telemetry by subscribing to specific xpaths of leaf nodes using Yang state models generated from ODL files. This subscription-based model provides precise and efficient data collection compared to the traditional RPC mechanism. Note that this solution supports only periodic streaming of telemetry data.

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

- **Stream telemetry data in gNMI-based message format over UDP (MX204, MX240, MX301, MX304, MX480, MX960, MX10003, MX10004, MX10008, MX10016, MX2008, MX2010, MX2020, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, and QFX5210)**—Junos OS uses a dial-out mechanism to send telemetry data to a collector over UDP. The message format is defined in the `jnx_gnmi_over_udp.proto` file. This mechanism supports only STREAM mode with SAMPLE as subscription mode. The message contains full key name and value pair information so the collector does not require data models for processing or consuming the telemetry data.

No Link Title, No Link Title, No Link Title, No Link Title, No Link Title, and [Junos YANG Data Model Explorer](#).]

MAC Learning

- **Support to configure system timezone (QFX Series)**—By default, mac-learning-logs use UTC timestamps. The logs appear in the output of the `show ethernet-switching mac-learning-log` command. Configure them to display in the system timezone for better alignment with local time. Use the `mac-learning-log system-timezone` option in the [l2-learning](#) command to display logs in the system timezone.

[See [l2-learning](#).]

Multicast

- **Hash-based PIM RPF selection (QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—Use hash-based Protocol Independent Multicast (PIM) reverse path forwarding (RPF) selection to route multicast traffic for a specific source and group (S, G) consistently through the same upstream node. The device uses multi-level hashing of PIM neighbor attributes such as router ID, cluster ID, and interface ID for consistent routing. For configuration, include the `hash-based-rpf-selection` statement under the `edit protocols pim` hierarchy.

[See [Hash-based PIM RPF selection](#), and [hash-based-rpf-selection](#).]

Network Address Translation (NAT)

- **Static NAT and PAT (QFX5120)**—Use static NAT and NATPT to translate private and public addresses and ports for transparent routing and address conservation. Define one-to-one mappings for source NAT and destination NAT to replace source or destination IP addresses between internal and external realms. For TCP and UDP, use static NAT and destination NAT to map both IP addresses

and transport ports so multiple hosts can share a single external address. The system updates IP and transport checksums automatically during translation.

[See [Configure Static NAT and PAT for QFX5120](#).]

Post-Quantum Cryptography (PQC)

- **PQC signatures for software images with off-box verification (ACX Series, EX Series, MX Series, QFX Series, and SRX Series)**—Use post-quantum cryptography (PQC) signatures to ensure software images remain unaltered, protecting integrity and guarding against quantum threats. The images comply with algorithms recommended by Commercial National Security Algorithm 2.0 (CNSA 2.0):
 - ML-DSA-87 PQC algorithm for digital signatures
 - SHA-512 for hashing

For off-box verification, request the PQC signature from the support portal. Confirm the image's SHA-512 hash, retrieve the public key from the certificate, and validate the signature with your chosen verifier. PQC signatures provide additional security beyond existing legacy signatures.

[See [PQC Signatures for Software Images](#).]

Routing Policy and Firewall Filters

- **New options for the show route command (QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, and QFX5210)**—We've introduced enhancements to the show route command to enable network operators understand and troubleshoot routing tables more easily. Use the prefix-length-distribution option to display counts of prefix lengths across routing tables for each instance. You can run the show route *destination* covering-subnets command to walk up the radix tree and list all routes covering a destination. These options provide deeper insight for troubleshooting and optimize routing behavior.

[See [show route](#).]

- **Apply firewall filters to loopback address logical interfaces on a per VRF basis (QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—By default, loopback firewall filters apply only to lo0.0, which is global and applicable to all control traffic entering the routing engine, regardless of VRF instances. Use the loopback-firewall-per-vrf configuration statement to apply firewall filters to loopback address logical interfaces on a per VRF basis.

[See [loopback-firewall-per-vrf](#).]

Routing Protocols

- **BGP support for deterministic path forwarding (DPF) in a CLOS network (EX4100-48MP, EX4100-24MP, EX4100-48P, EX4100-48T, EX4400-48MXP, EX4400-48P, EX4400-48T,**

QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)—BGP Deterministic Path Forwarding (DPF) divides a physical fabric into multiple logical fabrics so that different flows can be mapped to different logical fabrics to serve the requirements of the flows. Single hop External BGP (EBGP) might not meet all DC flow requirements, especially for drop and latency sensitive AI ML flows. BGP DPF colors the single-hop EBGP session on each link with a fabric color. If a link belongs to the red fabric, the EBGP session over the link is colored red. A route with no color community is advertised over any colored or uncolored EBGP sessions.

To configure fabric colors for BGP neighbors include the `fabric-color com-name` statement at the `[edit protocols bgp]` hierarchy level. You can also configure this feature at the BGP group or neighbor level.

To advertise colored IPv4 or IPv6 routes, include the `fabric-advertise` statement at the `[edit protocols bgp]` hierarchy level.

[See [BGP Deterministic Path Forwarding in a CLOS Network](#) .]

- **IS-IS multi-instance support on a single interface (ACX5448, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX10002-60C, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016)**—We have enhanced the IS-IS multi-instance feature to support multiple IS-IS instances on the same logical interface with instance identifier TLV 7.

Include the `instance-id` statement at the `[edit protocols isis-instance name]` hierarchy level.

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

- **Proactive IPv6 Neighbor Detection (QFX5110)**—Proactive IPv6 neighbor detection enables Junos OS to periodically discover and verify the reachability of IPv6 hosts connected to an interface. When enabled, the system sends Neighbor Solicitation messages for a configured IPv6 address range and updates the IPv6 neighbor cache upon receiving Neighbor Advertisements. You enable the feature globally and configure host discovery parameters per interface, including discovery interval and aging time out. The system refreshes learned neighbors when the aging timer expires and retries unresolved addresses based on the configured discovery interval. This feature helps operators proactively validate host reachability and maintain accurate neighbor cache entries in environments where host routes are not dynamically learned.
- **NDP Proxy Support with VRRP (SRX1500, SRX5400, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200 and QFX5210)**—NDP proxy support with VRRP enables a router to use the VRRP virtual MAC address when replying to IPv6 Neighbor Solicitation requests and DAD requests on behalf of other hosts. When you configure an interface with NDP proxy or DAD proxy and VRRP, only the VRRP master generates proxy Neighbor Advertisements. VRRP backup routers drop proxy requests and do not send proxy Neighbor Advertisement responses. This behavior supports both interface restricted and interface unrestricted proxy modes and maintains consistent address resolution during VRRP failover by keeping the MAC address stable. The feature also supports proxying for protocol installed IPv6 user host routes when user route proxy is enabled.

Additional Features

We've extended support for the following features to these platforms.

- **Supported transceivers, optical interfaces, and DAC cables.** (EX4100, EX4400, EX4400-24X, EX4400-EM-4Y, EX4650, MX304, MX10004, MX10008, QFX5120, and SRX4700). Select your product in the [Hardware Compatibility Tool](#) 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 optical transceiver becomes available.
- **Support for RFC 6395** (QFX5120-32C, QFX5120-48T, QFX5120-48Y, QFX5120-48YM, QFX5200, and QFX5210)

[See [cluster-id](#) and [disable-interface-id-tlv](#).]

- **Renaming OpenSSH implementation to JSSH (all platforms)**—The OpenSSH implementation in Junos OS is renamed "JSSH" to highlight its customized nature.

[See [ssh](#).]

What's Changed

IN THIS SECTION

- [Dynamic Host Configuration Protocol](#) | 128
- [EVPN](#) | 128
- [General Routing](#) | 128
- [Interfaces](#) | 131
- [MPLS](#) | 131
- [Network Management and Monitoring](#) | 131
- [Routing Protocols](#) | 131
- [Subscriber Access Management](#) | 132
- [User Interface and Configuration](#) | 132

Learn about what changed in this release for QFX Series Switches.

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

EVPN

- **EVPN-VXLAN T5 host-route auto-generated community (EX4100-24T, EX4300-MP, EX4400-24MP, EX4400-24P, EX4400-48F, EX4400-48MP, EX4650, MX960, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)**—We added `from interface` as an optional qualifier for a policy adding a Type 5 host-route auto-generated community to locally learned host routes.

[See [EVPN-VXLAN Pure Type 5 Host-Route Auto-Generated Community](#).] [PR1889160](#)

General Routing

- **Changes to `show system alarms` command output (QFX5130 and QFX5220)**—When the current version of the firmware is less than the minimum supported version, you can now see alarms for this mismatch in the output of the command. These alarms were not shown previously. For example, when you have a firmware version mismatch, you should now see output similar to the following:

```
user@host> show system alarms
18 alarms currently active Alarm time Class Description
2024-09-09 04:55:00 PDT Minor CHASSIS 0 BIOS ROM minimum supported firmware version mismatch
2024-09-09 04:55:20 PDT Minor CHASSIS 0 Fan CPLD minimum supported firmware version mismatch
2024-09-09 04:55:19 PDT Minor CHASSIS 0 Optics CPLD minimum supported firmware version
mismatch
```

[PR1836811](#)

- A new counter "Sessions hit due to high rate" is added to `show services service-sets screen-session-limit-counters` command for all subscriber traffic. This counter tracks the sessions that come up on the screen irrespective of the `alarm-without-drop` configuration. When `alarm-without-drop` option is disabled, all the counters display updated statistics. When `alarm-without-drop` is enabled, then, the screen-drop counters on `show services service-sets statistic screen-drop` command do not increase. The sessions hit due to high rate value is displayed.

[See [alarm-without-drop \(IDS Screen Next Gen Services\)](#), [show services service-sets statistic screen-drops \(Next Gen Services\)](#), and [show services service-sets statistic screen-session-limit-counters \(Next Gen Services\)](#).] [PR1849594](#)

- When you run the request vmhost zeroize command to zeroize a single Routing Engine on a dual Routing Engine device, the CLI incorrectly displays a message indicating that it will zeroize both Routing Engines. [PR1869854](#)
- **Validation of /vm-primary mount during JDM installation or upgrade (Junos Node Slicing External Server Deployment)**—During installation or upgrade of the Juniper Device Manager (JDM) on external servers running Junos Node Slicing Release 25.2 or later, an issue occurred with the validation of the /vm-primary mount that stores the GNF images. When /vm-primary was mounted using logical volumes (LVM), JDM would fail to detect that the underlying storage was an SSD. This issue is now fixed. However, the fix introduces a new dependency on the LVM2 package in the host OS. This package is included by default in standard installations of both RHEL and Ubuntu external servers. However, it is advised that you check if the LVM2 package is already installed on the host before installing or upgrading JDM. [PR1877593](#)
- On the MPC7E-10G line card, when you configure the 10-Gigabit Ethernet ports to operate as 1-Gigabit Ethernet ports, use the speed statement at both the edit interfaces interface name gigheter-options and edit interfaces interface name hierarchy levels. [PR1879198](#)
- **Control Maximum 802.1X Client Connections per Interface**—By default, dot1x interfaces configured in multiple supplicant mode have a client limit of 100 authenticated connections per interface. Any additional connection attempts beyond this limit will be automatically blocked. [PR1882148](#)
- **New option for debug collector data storage path** --We've included the option outdir to specify an output directory for storing debug collector data in a customised path. This allows you to organise and access diagnostic information more efficiently, adapting storage to your specific requirements.

[See [request system debug-info](#).] [PR1889710](#)
- **High-power optics support with CLI configuration option (QFX5240 and QFX5241)** —You can enable high-power optics across all ports by configuring the high-power-mode option for each port. This feature supports up to 32 high-power modules, allowing you to benefit from enhanced connection capabilities. Ensure you configure the necessary settings to initialize and utilize high-power optics effectively, optimizing your network's performance. [PR1891150](#)
- **Static flowspec prefix validation**—Commit validation now rejects static flowspec destination and source prefixes that use a shorter mask than the configured address family permits. The configuration must use a valid prefix length before the commit can succeed. This alignment with static route validation prevents invalid flowspec entries from being installed and ensures consistent routing policy enforcement across IPv4 and IPv6 configurations. [PR1903233](#)
- **Enhanced telemetry aggregation logic alignment with SNMP counter calculations (QFX Series)**—The Junos OS telemetry is updated to include missing internal counters used in SNMP (e.g., bucket drop,

idiscards, il3_incomplete, etc.). The change to the OpenConfig sensors ensures parity across both monitoring mechanisms.

[See [YANG Data Model Explorer](#).]PR1903308

- The `run test chassis fru` command now displays data for only the FRU (field replaceable units) that are supported on the QFX series and EX series devices and platforms.[PR1905481](#)
- **Enhanced ping command XML output (Junos OS and Junos OS Evolved)**—The XML output for the CE-IP ping has enhanced details to differentiate between successful and failed ping attempts. The update introduces new tags that track the total number of packets sent and received, while also highlighting host unreachable replies and the overall packet loss percentage.

[See [Supported and Unsupported Features for CE-IP Ping](#).] PR1906967

- **Sflow sampling with firewall filters (QFX5120, QFX5200, and QFX5210)**—This feature enables sFlow sampling through firewall filter actions. You can selectively sample traffic matching specific filter terms using the `then sampling` action, providing granular visibility while reducing sampling overhead. You can configure sampling as `set firewall family inet filter filter-name term term-name then sampling` and combine with standard sFlow configuration under `set protocols sflow hierarchy`, such as (sample-rate, collectors, interfaces) for configuring sFlow.

Note that when you configure the `sampling action` in a term, you cannot configure the `discard` and `reject action` in the same term.

[See [sFlow Support on Switches](#).]PR1910068

- **Hidden `disable-fpc` CLI statement (QFX5000 Series Switches)**—We've hidden the `disable-fpc` CLI statement at the following hierarchy levels:
 - `[edit system ddos-protection global]`
 - `[edit system ddos-protection protocols protocol-group aggregate]`

Earlier to this release, when you configure the `disable-fpc` CLI statement, the device generated DCPFE core when the device processes a large volume of data as it restricts the DDOS protection of the CPU.

[See [disable-fpc \(DDoS\)](#).]PR1910643

- Support added for new payload match conditions (EX4100 , EX4400 , EX5200, and QFX5120)- Support added for `payload-source-ipv4-address`, `payload-destination-ipv4-address`, `payload-source-ipv6-address`, `payload-destination-ipv6-address`, `payload-source-mac-address`, and `payload-destination-mac-address` match conditions for IPv4 and IPv6 traffic which can be applied on the ingress port only.

[See [Firewall Filter Match Conditions for IPv4 Traffic](#).]PR1914083

- **request jnu role delete Command Unhidden (Junos OS and Junos OS Evolved)**—The `request jnu role delete` command, which was previously hidden, is now available for use. This command allows users to delete the role of controllers or satellites directly through the CLI. This enhancement improves administrative control and flexibility in role management within the network infrastructure.

[See [request jnu role](#).]PR1930291

Interfaces

- **Improve heap memory utilization in link aggregation group (LAG) interfaces in large-scale configurations**—If you T5 host-route auto-generated community have a configuration with many interfaces in a large number of VLANs, and you've configured those interfaces as part of a logical aggregated Ethernet interface bundle, the device can approach 100% of its heap memory utilization. To avoid this problem in such large-scale environments, you can set the `lag-heap-optimization` option to reduce heap memory utilization on the device. When you configure this statement, the Packet Forward Engine on the device restarts.

[See [lag-heap-optimization](#).]

MPLS

- **Label field naming update in show route forwarding-table extensive output (Junos OS and Junos OS Evolved)**—The CLI output for MPLS forwarding entries now uses "Label" instead of "Load balance label." This name change standardizes the field name and improves overall clarity when reviewing MPLS forwarding information. The change affects only the output text and does not modify label allocation, forwarding behavior, or traffic handling. No configuration changes are required.PR1899543

Network Management and Monitoring

- **Shell Command Logging Enhancement**—All shell commands executed on the device and root sessions are now logged by default. This enhancement ensures enhanced security and auditability by capturing all commands entered in any shell environment, preventing bypassing audit logging.PR1867216
- **Deprecation of shell option**—The `shell` option no longer requires a separate configuration and is now the default behavior. Deprecating the shell option enhances efficiency and simplifies management tasks.

Routing Protocols

- **SRv6 MSD advertisement compliance**—MSD type 43 (SRH Max T.insert) is no longer advertised in IS-IS SRv6 node capability TLVs to match the latest standards draft. This alignment maintains

interoperability with peer devices and tools that expect the reduced MSD type list and ensures SRH limits are reported consistently across the network.[PR1879051](#)

- **Auto BGP peer discovery for unnumbered interfaces of iBGP peers**—Junos supports single-hop iBGP sessions for unnumbered peering. Previously, the iBGP implementation created an indirect next-hop for route resolution, even if the iBGP peer was directly connected.[PR1884861](#)
- **BGP route-refresh warning messages added**—New messages have been introduced to the `show bgp diagnostics warnings` command output. These messages provide diagnostic information when the number of BGP route-refresh messages exceed predetermined thresholds within specific time intervals of 5 minutes, 30 minutes, and 1 hour.

[See `show bgp diagnostics warnings`.][PR1909467](#)

- **BGP UPDATE messages must have the marker field set to all ones**—The marker field in BGP header is a 16-octet field included for compatibility that must be set to all ones. All Junos OS and Junos OS Evolved devices send a notification and tear down the BGP session if the marker field in the header of the received update message is not set to all ones.[PR1955678](#)

Subscriber Access Management

- **Addition of message-authenticator and no-message-authenticator attributes under [edit access radius-server], [edit access profile radius-server], and [edit system radius-server] hierarchies**—set `message-authenticator` if you require the RADIUS server to include the Message Authenticator attribute in replies to Access-Request messages. set `no-message-authenticator` to not require that attribute.[PR1871147](#)

User Interface and Configuration

- **Updated Annotate Command Behavior**—The `annotate` command now correctly handles multi-line comments that start with ``#``. This change ensures that all lines of a multi-line comment remain comments, preventing unauthorized configuration changes when the configuration is committed, reloaded, or rolled back. These improvements enhance security by mitigating potential privilege escalation attacks from users with limited configuration editing permissions.[PR1868636](#)
- **Enhanced Permission Checks for Rename/Copy Operations**—New permission checks have been introduced for rename and copy operations within the configuration hierarchy. These checks parse the hierarchy being modified to ensure the user has the required permissions for the hierarchy and its sub-hierarchies. If permissions are insufficient, the operation will fail and a "Permission denied" error will be displayed. This enhancement ensures that configuration modifications are performed only by authorized users, improving security and preventing unintended changes to critical system settings.[PR1882303](#)
- **Stale ui-state.db data in persistent NETCONF sessions post-mgd restart**—Existing NETCONF sessions might fetch stale data from `ui-state.db` after `mgd -N` restart. New sessions correctly map the

refreshed database. Scripts must establish new sessions post-restart to access updated values. Functional configuration remains unaffected. [Script failures monitoring "local-host" NETCONF sessions]-Scripts might fail when including "local-host" NETCONF sessions in monitoring operations. Internal sessions are now excluded from tracking. Scripts must filter out "local-host" sessions. No impact to internal application functionality.[PR1888557](#)

- **Generate genstate YANG modules on Junos devices**—You can use the `show system schema operational` command or equivalent RPC to generate the genstate YANG modules in the specified output directory on a device.

[See [show system schema](#).][PR1910388](#)

Known Limitations

IN THIS SECTION

- [Infrastructure](#) | 133

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

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

Infrastructure

- When upgrading older release to Junos OS release 21.2 and later, validation and upgrade might fail. The upgrade requires using the 'no-validate' option to complete successfully. <https://kb.juniper.net/TSB18251>[PR1568757](#)

Open Issues

IN THIS SECTION

- [General Routing](#) | 134

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

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

General Routing

- On rebooting Data Center Interconnect (DCI) Gateway device, while the device is coming up, multicast traffic drop is observed in the highly scaled config. The traffic drop is observed in the following condition. The rebooted device was the EVPN DF on I-ESI and after coming up, the DF election is triggered and the device is elected as EVPN DF on I-ESI. After coming up and elected as DF, the device builds the multicast routes afresh. This results in traffic drop. [PR1872219](#)
- On QFX10K8/16 , IFL memory is not freed on non-local interface of FPC during configuration changes (eg: IFL delete/deactivate) for L3IFL/L2IFL on AE/Scalar interfaces. Fix is present in common code and risky. It is a day one issue and the per IFL leak is minimal (0.00016% of total Kernel heap size) .[PR1884163](#)

Routing Protocols

- On QFX3500 or QFX5100 switches, when parity errors occur on interfaces, they might affect the memory management unit (MMU) memories. MMU counters can be corrupted, the interface buffers might be stuck, and there might be interface flaps and traffic loss on the affected ports. As a workaround (restoration only), reboot the system. [PR1169700](#)

Resolved Issues

IN THIS SECTION

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

- JUNOS_REG: QFX51200-48YM: Fan status output was not same after/before device vc-switch over. [PR1758400](#)
- [QFX5120-32C] LEDs do not work properly with QSFP or without optics [PR1855372](#)
- The dcpfe process crashes on specific Junos QFX and EX platforms due to memory corruption [PR1856424](#)
- "cpu-utilization-idle" and "temperature-cpu" properties for FPC are not present in QFX Models. [PR1864591](#)
- The PTP packets are dropped when IGMP snooping is enabled [PR1873129](#)
- JMA package fails to initialise after a power cycle on EX4650/QFX-5E series devices [PR1882472](#)
- Error messages "LBCM-L2,brcm_port_family_detach(),8814:IFF detach failed for IFD esi IFL xxxxxx with err -1" will be observed on Junos QFX5k and EX4k Platforms [PR1883866](#)
- Next-hop entries are not getting programmed in ECMP unilist group after device upgrade [PR1886612](#)
- snmp mib walk DomCurrentLaneAlarms does not show proper lanes Values in the latest builds [PR1886889](#)
- Stale MAC routes pointing to SVLBNH cause traffic drop on QFX5K [PR1887225](#)
- QFX5120 - SFP+ Modules disappear/down post upgrade [PR1890867](#)
- PFE crash observed during VXLAN classifier unbind or EZ-LAG commit operations [PR1894833](#)
- Traffic loss will be observed when VPLAG is configured on Junos QFX5k and EX4k platforms [PR1895903](#)
- Flooding back QFX5120-48y-8c BUM traffic towards source. [1897459](#)
- Enabling Sflow action on TD3 [PR1899488](#)
- Significant system slowness is observed post software upgrade on QFX5120-48YM [PR1902869](#)
- Selection of VGA-IP as source IP for RE-ARP packet causes incorrect ARP updation of VGA-IP getting bound with IRB-MAC at end-hosts [PR1909786](#)

- Device crash when executing "show topology address" with invalid Input from VTY mode [PR1909975](#)
- EVPN routes are stuck in the KRT queue [PR1913519](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 148

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

Upgrading Software on QFX Series Switches

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

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to <https://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.
3. Select 25.4 in the Release pull-down list to the right of the Software tab on the Download Software page.
4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 20.3 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new jinstall package on the device.



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

Customers in the United States and Canada use the following command:

```
user@host> request system software add source/jinstall-host-qfx-5-x86-64-25.4-R1.n-secure-signed.tgz reboot
```

Replace *source* with one of the following values:

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

Adding the `reboot` command reboots the switch after the upgrade is installed. When the reboot is complete, the switch displays the login prompt. The loading process can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 20.3 jinstall package, you can issue the `request system software rollback` command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a **junos-vmhost-install-x.tgz**.

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot. If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.



NOTE: The QFX10002-60C switch supports only the 64-bit version of Junos OS.



NOTE: If you have important files in directories other than /config and /var, copy the files to a secure location before upgrading. The files under /config and /var (except /var/etc) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

If the installation package resides locally on the switch, execute the **request vmhost software add <pathname> <source>** command.

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-25.4R1.9.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-x86-64-25.4R1.9.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-x86-64-25.4R1.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-x86-64-25.4R1.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-x86-64-25.4R1.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-
x86-64-25.4R1.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.

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 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 11: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
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 SRX Series

IN THIS SECTION

- [What's New | 149](#)
- [What's Changed | 158](#)
- [Known Limitations | 165](#)
- [Open Issues | 166](#)
- [Resolved Issues | 166](#)
- [Migration, Upgrade, and Downgrade Instructions | 170](#)

What's New

IN THIS SECTION

- [Application Identification \(AppID\) | 150](#)
- [Application Layer Gateways \(ALGs\) | 150](#)
- [Chassis | 150](#)
- [Device Security | 151](#)
- [High Availability | 151](#)
- [Identity Aware Firewall | 152](#)
- [Interfaces | 153](#)
- [Juniper Advanced Threat Prevention Cloud \(ATP Cloud\) | 153](#)
- [Junos OS API and Scripting | 153](#)
- [Junos Traffic Vision | 154](#)
- [MACsec | 155](#)
- [MPLS | 155](#)
- [Network Management and Monitoring | 156](#)
- [Post-Quantum Cryptography \(PQC\) | 156](#)

- [Public Key Infrastructure \(PKI\) | 157](#)
- [Additional Features | 157](#)

Learn about new features introduced in this release for SRX Series devices.

Application Identification (AppID)

- **Authentication support in proxy profiles (SRX Series, cSRX, and vSRX)**—You can enhance your network's security by configuring authentication within proxy profiles, allowing secure access to external feeds and services. By setting a username and password, you enable authenticated HTTPS communication through a proxy, supporting services such as SecIntel, AppID, Content Security, and PKID. This configuration prevents unverified data sources from accessing protected environments. To configure the feature, set a username and password in your proxy profile and use HTTPS to encrypt credentials, mitigating security vulnerabilities. Regularly update passwords and monitor for unauthorized access to maintain security integrity.

[See [Configuring SSL Proxy](#).]

Application Layer Gateways (ALGs)

- **Support for client identifier in forwarded DNS queries using experimental EDNS(0) option (NFX150, NFX250, NFX350, cSRX, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—You can enhance DNS query responses by utilizing the *Client ID in Forwarded DNS Queries* feature on SRX Series devices. This functionality supports precise client identification, essential for services such as parental control. You can configure SRX Series Firewalls to include client identifiers such as MAC addresses, IPv4, or IPv6 addresses in DNS queries. This feature, intended for controlled environments, generates targeted DNS responses based on the originating device's identity, improving service accuracy and network efficiency.

[See [DNS ALG](#).]

Chassis

- **Improvement to resiliency features (SRX1600, SRX2300, SRX4120, and SRX4300)**—We've made multiple improvements to resiliency features on the listed devices, including:
 - Support for a minimum of three failed readings before triggering an alarm to improve reliability
 - Alarms for voltage output, current, and fan faults for AC and DC power supply units (PSUs)
 - Range check for all sensors to ensure that only valid sensor values can trigger alarms

- Availability of debug information to triage root cause for alarms, in addition to monitoring component health, triggering alarms and log messages. Recovery is not supported.
- Alarm for mismatching firmware when the firmware installed does not follow the minimum requirements and recommendations

[See [Resiliency](#).]

Device Security

- **VRF-aware zone-based security policies with flow enhancements for Layer 3 VPN over EVPN-VXLAN and MPLS (SRX Series and vSRX 3.0)**—You can enforce security policies for each virtual routing and forwarding (VRF) instance by creating VRF-aware security zones. This approach helps improve scalability and intuitive policy management across EVPN-VXLAN and Layer 3 VPN (L3VPN) segments. You can define zones by VRF instance, not by VRF group, and implement intra-VRF or inter-VRF policies using the CLI or the management user interface.

[See [Security Policies with VRF-Aware Security Zones](#), [Flow Management in SRX Series Devices Using VRF Routing Instance](#), [show security flow session](#), and [security-zone](#).]

- **Group-based policy in VXLAN architecture (SRX Series and vSRX)**—Use a group-based policy (GBP) to create microsegmentation in VXLAN architecture by defining application-centric policies. Associate endpoints with tags that identify business functions to manage network access and direct traffic between endpoint groups. Enforce granular access control using new match options for source and destination tags. This approach strengthens security and simplifies policy enforcement across campus networks.

[See [Group-Based Policies in VXLAN Environments](#).]

- **FQDN ID for enhanced policy management and dynamic IP resolution (SRX Series, cSRX, and vSRX)**—Use unique fully qualified domain name (FQDN) ID mappings to manage frequent IP address changes. The system stores each FQDN's identifier (ID) in the Routing Engine and Packet Forwarding Engine. This FQDN ID storage enables quick lookups without constant policy updates, improving stability. This feature runs by default.

[See [DNS Snooping for Security Policies](#).]

High Availability

- **Selective session synchronization for MNHA (SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—Use selective session synchronization for Multinode High Availability (MNHA) during cold and hot synchronization to optimize performance, reduce redundant state replication, and maintain fine-grained control over synchronization durations. Disable synchronization for short-lived traffic or defer synchronization with a minimum age. You can configure selective session synchronization using default or user-

defined flow profiles. You can disable session synchronization for short-lived sessions by setting `session-sync disabled` or delay synchronization based on session age by adjusting `session-sync-min-age`.

[See [Selective Session Synchronization for Multinode High Availability](#).]

- **Four-node MNHA (SRX4600 and SRX4700)**—Use four-node Multinode High Availability (MNHA) to strengthen continuity by deploying two MNHA pairs across domains, including separate data centers. Each pair uses an interchassis link (ICL), and the pairs interconnect with an interdomain link (IDL) for secure intra-domain communication and failover if one pair becomes unavailable. The design supports SRG0 services such as firewall and NAT but does not support SRG1+ services—for example, IPsec VPN. Four-node MNHA provides resilience against localized disruptions. Four-node MNHA support is available only for the routing mode (Layer 3) of MNHA.

[See [Four-Node Multinode High Availability](#).]

- **IDL HA link encryption (SRX4600 and SRX4700)**—You can extend high availability across data center domains with four-node Multinode High Availability (MNHA). An interdomain link (IDL) synchronizes control and data plane states between nodes across the domains. You can secure IDL traffic using IPsec with IKEv2, multiple security associations (SA), AES-GCM-256 encryption, and use either preshared keys (PSKs) or public key infrastructure (PKI) for authentication. To encrypt the HA link for the IDL, install the Junos OS IKE package on SRX Series Firewalls and configure a VPN profile for HA traffic. Include the `ha-link-encryption` in your IPsec VPN configuration. An encrypted IDL link ensures secure interdomain communications.

[See [Four-Node Multinode High Availability](#).]

Identity Aware Firewall

- **Optimization of IPC messages for unified access control (UAC) authentication entries (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Optimize system performance by using IPC message communication between the Routing Engine and Packet Forwarding Engine. Group multiple role entries into one message to improve efficiency. Use the `clear services unified-access-control authentication-table` command to refresh the UAC authentication table for accurate role and user data.

[See [clear services unified-access-control authentication-table](#).]

- **User identity through HTTP XFF header (SRX Series Firewalls, and vSRX3.0)**—You can identify users behind proxies by extracting the originating client IP from HTTP X-Forwarded-For (XFF) or Forwarded headers. Avoid relying on the packet source IP. Use this method to prevent misattribution from proxy addresses and to improve policy enforcement, logging, and analytics accuracy. Use the `set services user-identification forward-header-lookup` command to enable the XFF header feature.

[See [forward-header-lookup](#), [show services user-identification forward-header-statistics](#), and [clear services user-identification forward-header-statistics](#).]

Interfaces

- **Logical interfaces limit expansion (SRX1600, SRX4120, SRX2300, and SRX4300)**—The device raises the logical interface limit from 4000 to 8000 for large-scale deployments requiring more VLANs or tunnel interfaces. You cannot configure 8000 IFLs on one revenue port because the IEEE 802.1Q VLAN ID field uses 12 bits, limiting VLANs to 4095. Overcome this limitation by using two or more revenue ports to configure more than 4000 IFLs. Review platform prerequisites and operational impacts.
- **QSFP breakout channelization for 4x10G on PIC port 0 with fixed subports 1–4 (SRX4700)**—Use QSFP breakout optics to enable 10 G breakout on PIC port 0 by setting 4x10G channelization in port-profile config-D. Only port 0 supports 4x10 Gbps on a 1x400 Gbps/4x100 Gbps/8x50 Gbps PIC, with a fixed range of one through four subports. Configure port 0 to 4x10 Gbps, 3x10 Gbps, 2x10 Gbps, or 1x10 Gbps. Use the `set chassis fpc <fpc slot> pic <pic slot> port <port number> channel-speed <10G | 25G>` command to configure channelization. Commit changes with any configuration or deletion of unused settings on port one to ensure deterministic behavior. [See [Port Speed on SRX4700 Firewalls](#).]

Juniper Advanced Threat Prevention Cloud (ATP Cloud)

Flow-based AV signature exempt lists (SRX Series Firewalls and vSRX)—Use flow-based antivirus (AV) signature exempt lists to exclude specific signatures from malware inspection. Configure these lists either through the Juniper ATP Cloud portal or by using CLI commands. Use CLI commands to add, delete, export, or import file signatures. Use signature exempt lists to reduce false positives and ensure more accurate threat detection.

[See [Create Allowlists and Blocklists](#).]

Junos OS API and Scripting

- **Support for libslax 3.1.6 and SLAX version 1.3 (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—We've upgraded the libslax library to version 3.1.6, which corresponds to SLAX version 1.3. The upgraded library includes:

- Slax processor enhancements including a new mode, additional options, and simplified argument parsing
- New libslax extension library functions
- Improved SLAX syntax options
- New SLAX functions and enhancements to existing functions and statements
- Hexadecimal numbers support in SLAX scripts

This update aligns SLAX processing with contemporary scripting conventions, ensuring efficient, readable scripting while maintaining backward compatibility.

[See [libslax Distribution Overview](#).]

Junos Traffic Vision

- **Junos Traffic Vision (SRX Series Firewall and vSRX 3.0)**—We now support Junos Traffic Vision (previously known as Jflow) on SRX Series Firewall. Junos Traffic Vision uses security policies and logical systems to generate flow records that capture addresses, packet counts, and byte counts. NAT44, NAT64, and NAT66 session template records include new information using J-Flow version 9 and IPFIX.

The new elements are:

IANA IPFIX ID	Field Name	Size (Bytes)
152	flowStartMilliseconds	32
6	tcp_flags	16
10	input_snmp	32
14	output_snmp	32
95	application_id	32
234	ingressVRFID	32
235	egressVRFID	32

Specify only the required template events and assign a predefined value (256 through 65535) for reuse.

Templates also support non-NAT firewall IPv4 and IPv6.

[See [Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250](#) and [Junos Traffic Vision Support on MS-MIC and MS-MPC](#).]

MACsec

- **Support for MACsec features over WAN (SRX1600, SRX4120, and SRX2300)**—Configure Media Access Control Security (MACsec) on logical interfaces to extend the benefits of hop-to-hop MACsec security to point-to-point security. Use MACsec features on logical interfaces to establish more secure VLAN-level MACsec connections in enterprise WAN and service provider networks. When these devices are in routing mode, they support:
 - Custom EAPoL destination MAC address for unicast MAC multicast, PAE, provider bridge, and LLDP multicast
 - MACsec on logical interfaces for Layer 2 or Layer 3 with VLAN tagging
 - Single-tagged VLAN IDs in clear text to support VLAN-level MACsec
 - GCM-AES-128, GCM-AES-256, GCM-AES-XPB-128, GCM-AES-XPB-256 cipher suites
 - Unencrypted MACsec
 - Static CAK security mode
 - MACsec using pre-shared key (PSK) hitless rollover keychain
 - Boundary delay
 - 802.1X authentication (dot1x protocol) for improved security
 - Fail open mode (should-secure) and must secure mode (default). The configurations for must-secure and should-secure are mutually exclusive. Only configure one option on a given physical interface for MACsec logical interface sessions. However, you can configure different options on different physical interfaces.

Before configuring these features, ensure there is Layer 2 adjacency between the customer edge devices. Then, enable MACsec on a logical interface using the unit *unit-number* option at the [edit security macsec interface *interface-name*] hierarchy level.

[See [Configuring MACsec, Media Access Control Security \(MACsec\) over WAN](#), and [Configuring Advanced MACsec Features](#).]

MPLS

- **MBGP MPLS L3VPN with packet mode and high-availability interfaces (SRX4600 and SRX4700)**—Use Multiprotocol BGP (MBGP) MPLS L3VPN on SRX4600 and SRX4700 with ASIC-based hash

forwarding for high performance. SRX4600 and SRX4700 sessionize and MPLS-encapsulate packets from edge devices. SRX4600 and SRX4700 parse packets from provider routers, match the packets to existing or new sessions, and then decapsulate the packets without installing sessions into the ASIC. Configure MPLS in packet mode to use these platforms as provider devices.

[See [MPLS Applications User Guide](#), [Layer 3 VPNs User Guide for Routing Devices](#), and [BGP User Guide](#)].

Network Management and Monitoring

- **Logging infrastructure support for transport-level statistics and stream-based counters (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX4100, SRX4120, SRX2300, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)**— We've introduced a logging infrastructure to manage transport-level statistics for logs sent over TCP, UDP, and SSL protocols. This transport-level logging infrastructure provides a comprehensive framework to monitor log delivery performance across these transport protocols.

Using the CLI commands, users can gain insights into log delivery performance by tracking metrics such as:

- Bytes sent and dropped
- Connection details
- Delivery errors

The logging infrastructure enables users to configure up to eight security log streams, each associated with a specific host IP address. This approach provides detailed analysis of stream-specific statistics and counters.

Use the following commands to view and clear the log statistics across transport protocols for multiple streams:

- `show security log transport`
- `clear security log transport.`

[See [show security log](#), [show security log transport](#), [clear security log](#), and [clear security log transport](#).]

Post-Quantum Cryptography (PQC)

- **PQC signatures for software images with off-box verification (ACX Series, EX Series, MX Series, QFX Series, and SRX Series)**—Use post-quantum cryptography (PQC) signatures to ensure software images remain unaltered, protecting integrity and guarding against quantum threats. The images comply with algorithms recommended by Commercial National Security Algorithm 2.0 (CNSA 2.0):
 - ML-DSA-87 PQC algorithm for digital signatures

- SHA-512 for hashing

For off-box verification, request the PQC signature from the support portal. Confirm the image's SHA-512 hash, retrieve the public key from the certificate, and validate the signature with your chosen verifier. PQC signatures provide additional security beyond existing legacy signatures.

[See [PQC Signatures for Software Images](#).]

Public Key Infrastructure (PKI)

- **HTTPS support for PKI (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—With HTTPS support for PKI, you can enhance the security of certificate management operations. This feature establishes secure communication channels for SCEP enrollment and CRL revocation, protecting sensitive information. The PKI process dynamically selects HTTP or HTTPS based on configured URLs, providing flexibility and secure transmissions.

[See [PKI Components in Junos OS](#).]

Additional Features

We've extended support for the following features to these platforms.

- **Supported transceivers, optical interfaces, and DAC cables.** (EX4100, EX4400, EX4400-24X, EX4400-EM-4Y, EX4650, MX304, MX10004, MX10008, QFX5120, and SRX4700). Select your product in the [Hardware Compatibility Tool](#) 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 optical transceiver becomes available.
- **Support for NAT IPv6 translations offloading to NPU and DS-Lite in SOF (SRX4700)**

[See [IPv6 NAT Overview](#) and [IPv6 Dual-Stack Lite](#).]

- **Support for file system encryption with TPM 2.0 (SRX1600, SRX2300, SRX4120, SRX4300, and SRX4700)**

[See [File System Encryption](#).]

- **Renaming OpenSSH implementation to JSSH (all platforms)**—The OpenSSH implementation in Junos OS is renamed "JSSH" to highlight its customized nature.

[See [ssh](#).]

What's Changed

IN THIS SECTION

- Application Security | 158
- Authentication and Access Control | 159
- Chassis Clustering | 159
- Content Security | 159
- Dynamic Host Configuration Protocol | 160
- Flow-Based and Packet-Based Processing | 160
- High Availability | 161
- Identity Aware Firewall | 161
- Infrastructure | 161
- Interfaces | 161
- Intrusion Detection and Prevention (IDP) | 161
- Network Management and Monitoring | 162
- Platform and Infrastructure | 162
- Routing Policy and Firewall Filters | 163
- Routing Protocols | 163
- SSL Proxy | 163
- User Interface and Configuration | 163
- VPNs | 163

Learn about what changed in this release for SRX Series.

Application Security

- **Support for routing-instance and source address for application signature download (SRX Series Firewalls)**—New configuration options enable you to specify a custom routing instance and source address for downloading application identification signature packages. This change enables enhanced traffic routing control and aligns signature downloads with specific network policies.

[See [Configuring Application Signature Package Download Options.](#)]

Authentication and Access Control

- For push-to-identity-management to successfully push the authentication entry to JIMS, you must configure JIMS and verify that JIMS status is online.

[See [push-to-identity-management](#).]

Chassis Clustering

- **MNHA Upgrade Requirement for IPv6-Based ICL Encryption (SRX4200, SRX4300, SRX4600, and SRX4700)**—In Multinode High Availability (MNHA) deployments that use IPv6 addresses for Interchassis Link (ICL) (HA link) encryption, upgrading from an earlier Junos OS release to a release that supports IPv6-based ICL encryption requires following the isolated node upgrade procedure to ensure a successful transition. This is a one-time requirement when moving to the first supported release; subsequent upgrades do not require the isolated node upgrade procedure.

For isolated node upgrade procedure, see [\[SRX\] Isolated Node Upgrade Procedure for Multinode High Availability](#).

For Software Upgrade in Multinode High Availability, see [Software Upgrade in Multinode High Availability](#).

Content Security

- **Support for IPv6 in Enhanced Web filtering (EWF) (SRX Series Firewall)**—Use this feature to enable IPv6 support in server settings for Content Security Enhanced Web Filtering and Sophos Antivirus. Configure IPv6 addresses in proxy profiles and host name fields to improve compatibility and reachability across IPv6 networks. By default, IPv4 web filtering remains enabled.
- **Antivirus syslog enhancement (SRX Series Firewall)**—We've added the following fields in the antivirus syslog:
 - application
 - application-category
 - file-size
 - file-hash
 - malware-info
 - nested-application
 - policy-name
 - threat-score

[See [show log](#).]

- **Web filtering syslog enhancement (SRX Series Firewall)**—We've added the policy-name field in the Web filtering syslog.

[See [show log](#).]

- **Web filtering cache-preload status enhancement (SRX Series Firewall and vSRX)**—We've added the Download status field in the show security utm web-filtering cache-preload status operational command output.

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

Dynamic Host Configuration Protocol

- The show dhcp-security vlans detail command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

Flow-Based and Packet-Based Processing

- **Standardized byte accounting for non-VLAN Ethernet packets across SPC3 and IOC (SRX Series Firewall)**—Starting from this release, for an Ethernet header without VLAN, the byte count increment for packets sized 100 bytes has been standardized across SPC3 and IOC processing units.

Previously, SPC3 processed packets showed an increment of 86 bytes, while IOC packets showed 100 bytes. Now, both SPC3 and IOC units uniformly skip counting the Ethernet header, resulting in a consistent byte count increment of 86 bytes per packet. This enhancement ensures uniformity in packet processing and accurate byte count accounting across different processing units.

- **Enable or disable IEEE 802.3x pause for NICs in FPC 0 (SRX4700)**—SRX4700 allows you to control Lossless Flow Control (IEEE 802.3x Pause) on NICs associated with FPC 0 to optimize performance under burst-heavy or asymmetric CPU load conditions. By default, lossless flow control is enabled to provide maximum buffering and protect against congestion. You can disable flow control to prevent pause-frame propagation and ensure that traffic continues to reach underutilized CPUs even when other CPUs are congested.

Configure using:

- set chassis fpc 0 lossless-flow-control lossless-flow-control-enable
- set chassis fpc 0 lossless-flow-control no-lossless-flow-control-enable

This setting applies only to **NICs on FPC 0 (SRXPFE)**.

The configuration takes effect immediately (no reboot required) and persists across system **reboots**, **VMHOST reboots**, and **SRXPFE restarts**.

[See [Flow Control for Ethernet Interfaces](#).]

High Availability

- **IPv6 support for Multinode High Availability Configuration (SRX Series Firewalls)**—A new enhancement has been added to support the configuration of IPv6 addresses for the active signal route, backup signal route, and install on failure route options under services-redundancy-group configurations on your MNHA setup. With this update, you can now configure IPv6 addresses, facilitating compatibility with IPv6 networks and improving overall network interoperability.

[See [Multinode High Availability](#)]

Identity Aware Firewall

- **Identity Aware Firewall**—Firewall users must keep the captive portal web login page open after they successfully authenticate. The system automatically logs the user out of the captive portal when the login page is closed.

[See [Captive Portal Authentication](#) and [Configure Authentication Methods for SRX Firewall Users](#).]

Infrastructure

- You can now boot vSRX 3.0 with either UEFI or BIOS.

Interfaces

- **ARP restriction for VLAN IDs 3072 to 4094 (SRX4700)**—You cannot configure VLAN IDs ranging from 3072 to 4094. This ensures correct network behavior and prevents potential conflicts within these VLAN ranges, promoting network stability and reliability.

Intrusion Detection and Prevention (IDP)

- **Improved Handling of IDP Policy Compilation Status (SRX Series Firewall)**—Previously, if an IDP policy compilation failed and a subsequent commit did not involve IDP changes, the compilation status could be lost or appear blank. This has been resolved—the system now retains and displays the last known policy compilation status, even when later commits do not trigger policy recompilation or when the policy is unloaded due to configuration changes. There is no change in the underlying IDP functionality, only in how the status message is preserved.

Network Management and Monitoring

- **Deprecation of shell option**—The `shell` option no longer requires a separate configuration and is now the default behavior. Deprecating the shell option enhances efficiency and simplifies management tasks.
- **New option for debug collector data storage path**—We've included the option `outdir` to specify an output directory for storing debug collector data in a customised path. This allows you to organise and access diagnostic information more efficiently, adapting storage to your specific requirements.

[See [request system debug-info](#).]

- **Enhanced syslog fields for screen event monitoring (SRX Series Firewalls)**—We have added new fields to screen-related syslogs to improve network security monitoring and analysis. The syslogs now include additional fields such as source port, destination port, destination zone name, session ID, policy name, application user, threat score, and threat severity. These fields apply to various screen types like `RT_SCREEN_IP`, `RT_SCREEN_ICMP`, `RT_SCREEN_TCP`, `RT_SCREEN_TCP_DST_IP`, `RT_SCREEN_UDP`, and `RT_SCREEN_TCP_SRC_IP`. This enhancement allows for more detailed analysis and monitoring, improving threat detection and response capabilities.
- **IPv6 DNS resolution option in security log stream configuration (SRX Series Firewalls and vSRX3.0)**—You can enable the `prefer-ipv6-dns` option under the `show security log stream s1` host configuration hierarchy to prioritize IPv6 address resolution for DNS queries. This option ensures that IPv6 addresses are used instead of the default IPv4 addresses. This configuration enhances IPv6 network compatibility and supports environments that require IPv6 addressing.

Platform and Infrastructure

- **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](#).]

- You can now enable zeroization on a vSRX 3.0 Virtual Firewall using CLI to destroy Critical Security Parameters (CSPs). Run the `request system zeroize` command to zeroize the system configuration and keys. When you run this command all the configuration information is removed, and the key values are reset and the vSRX 3.0 firewall is reverted to factory defaults after reboot.

Routing Policy and Firewall Filters

- **IPv6 address range support in address book configuration**—You can configure IPv6 address ranges within the address book, enabling more flexible network management. With this feature IPv6 range configurations can be split into multiple prefixes. You must handle this feature carefully as it transforms ranges into multiple prefixes.

Routing Protocols

- **BGP UPDATE messages must have the marker field set to all ones**—The marker field in BGP header is a 16-octet field included for compatibility that must be set to all ones. All Junos OS and Junos OS Evolved devices send a notification and tear down the BGP session if the marker field in the header of the received update message is not set to all ones.[PR1955678](#)

SSL Proxy

- **Configuration Limits for SSL Proxy Profiles**—We have updated the limits for Trusted CA certificates, Server certificates, and URL categories in both SSL forward proxy and SSL reverse proxy configurations. These changes ensure compliance with the maximum configuration blob size limit of 56,986 bytes.
 - Trusted CA certificate/Server certificates: Maximum limit 400 (reduced from 1024)
 - URL categories: Maximum limit 800 (unchanged)

[See [Configuring SSL Proxy](#).]

User Interface and Configuration

- **Generate genstate YANG modules on Junos devices**—You can use `show system schema operational` command or equivalent RPC to generate the genstate YANG modules in the specified output directory on a device.

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

VPNs

- **Default installation of junos-ike package on additional platforms (SRX1500, SRX4100, SRX4200, SRX4600, and vSRX 3.0)**—The `junos-ike` package is installed by default on SRX1500, SRX4100, SRX4200, SRX4600, and vSRX3.0 firewalls, ensuring the default support for `iked` process for IPsec VPN service. This aligns with the existing default installation of the package on SRX5000 line with Routing Engine 3 (SRX5K-SPC3 with RE3). You can delete the `junos-ike` package using the command

request system software delete junos-ike. This runs the kmd process on these firewalls, allowing flexible management of your security infrastructure.

[See [IPsec VPN Overview](#).]

- **Global option to disable inline IPsec hardware offloading (SRX4700)**—You can disable hardware offloading of IPsec tunnel processing in the Packet Forwarding Engine ASIC. Use the command `set security ipsec hw-offload-disable` to globally disable this inline IPsec processing of packets. When you configure the statement, the firewall processes IPsec tunnels in CPU instead of the Packet Forwarding Engine ASIC. This statement replaces the previous hidden option `no-hw-offload` at the `[edit security ipsec]` hierarchy level. This global configuration provides a streamlined approach to managing IPsec hardware offloading settings at the firewall level.

[See [ipsec \(Security\)](#).]

- **Deprecation of weak algorithms in IPsec VPN (SRX Series and vSRX 3.0)**—We've deprecated the weak algorithms in IKE and IPsec proposals. You'll no longer be able to use the following algorithms:

Table 12: Deprecated Junos CLI Options

Type	Algorithm	Junos CLI Statement
Encryption Algorithm in IKE Proposal	des-cbc and 3des-cbc	<code>set security ike proposal <i>name</i> encryption-algorithm</code>
Authentication Algorithm in IKE Proposal	md5 and sha1	<code>set security ike proposal <i>name</i> authentication-algorithm</code>
DH Group in IKE Proposal	group1, group2, and group5	<code>set security ike proposal <i>name</i> dh-group</code>
Encryption Algorithm in IPsec Proposal	des-cbc and 3des-cbc	<code>set security ipsec proposal <i>name</i> encryption-algorithm</code>
Authentication Algorithm in IPsec Proposal	hmac-md5-96 and hmac-sha1-96	<code>set security ipsec proposal <i>name</i> authentication-algorithm</code>

You will receive a warning message if you configure these deprecated algorithms explicitly. As an alternative, we recommend that you configure the stronger algorithms to enhance the security in IPsec VPN.

[See [proposal \(Security IKE\)](#), and [proposal \(Security IPsec\)](#).]

- **SCEP certificate re-enrollment (SRX Series)**—The RFC8894 states that the challenge password is optional when an existing certificate signs a renewal request. The challenge password is not mandatory. You can commit the configuration without the challenge password.

[See [Enroll a Certificate](#).]

- **Configuration validation for HA link encryption (SRX Series)**—New validation checks have been introduced to restrict the configuration of tunnel MTU for HA link encryption tunnels in a Multinode High Availability setup. The validation check ensures that the end-to-end MTU for HA links using IPv6 encryption meets the minimum requirement of 2000 bytes, helping maintain optimal performance and reliability during high availability operations. For example, if your configuration includes the following stanza, you'll receive a commit check error: `user@host# set security ipsec vpn L3HA_IPSEC_VPN tunnel-mtu <bytes>`.



NOTE: In an MNHA setup, for IPv6 HA link encryption, ensure to maintain a minimum end-to-end MTU of 2000 bytes.

[See [Multinode High Availability](#).]

- **Support for hmac-sha-384/512 authentication in PMI (SRX Series Firewalls and vSRX 3.0)**—You can configure hmac-sha-384 and hmac-sha-512 authentication algorithms with PowerMode IPsec (PMI) when running IPsec VPN with the `iked` process.

[See [PowerMode IPsec](#).]

Known Limitations

IN THIS SECTION

- [Infrastructure](#) | 166

Learn about known limitations in this release for SRX Series devices.

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

Infrastructure

- When upgrading from Junos OS Releases before Junos OS Release 21.2 to Junos OS Release 21.2 and onward, validation and upgrade might fail. The upgrade requires using the 'no-validate' option to complete successfully. <https://kb.juniper.net/TSB18251>. [PR1568757](#)

Open Issues

IN THIS SECTION

- [Platform and Infrastructure | 166](#)

Learn about open issues in this release for SRX Series devices.

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

Platform and Infrastructure

- An Authentication Bypass by Spoofing vulnerability in the RADIUS protocol of Juniper Networks Junos OS and Junos OS Evolved platforms allows an on-path attacker between a RADIUS server and a RADIUS client to bypass authentication when RADIUS authentication is in use. Please refer to <https://supportportal.juniper.net/JSA88210> for more information. [PR1850776](#)
- Multiple vulnerabilities have been resolved in MQTT included with Junos by fixing vulnerabilities found during external security research. Please refer to <https://supportportal.juniper.net/JSA71655> for more information. [PR1651519](#)

Resolved Issues

IN THIS SECTION

- [Chassis Clustering | 167](#)

- [Flow-Based and Packet-Based Processing](#) | 167
- [General Routing](#) | 167
- [Infrastructure](#) | 168
- [Intrusion Detection and Prevention \(IDP\)](#) | 169
- [J-Web](#) | 169
- [Platform and Infrastructure](#) | 169
- [Routing Policy and Firewall Filters](#) | 169
- [Content Security](#) | 169
- [VPNs](#) | 169

Learn about the issues fixed in this release for SRX Series devices.

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

Chassis Clustering

- ISSU might not work and process stops on primary node during high availability upgrade. [PR1895134](#)
- Backup nodes stuck in cold synchronize failure after all FPCs reset due to SPC stops responding files in SRX Series Firewall chassis cluster. [PR1895790](#)

Flow-Based and Packet-Based Processing

- SRX Series Firewall drops to-the-box ICMPv6 echo request with sequence number 3503 and 35000 through 35999. [PR1892890](#)

General Routing

- MNHA Conn State and ICL are down after 48+ hours of device being up with background traffic. [PR1822662](#)
- SPC3 flowd heartbeat lost during interface failover. [PR1837905](#)
- SRX Series Firewall unstable NTP issue. [PR1859188](#)
- The flowd process might stop on all SRX Series Firewall in multicast scenario with PIM. [PR1877771](#)

- Disabled port becomes up after rebooting Junos SRX340, SRX300, SRX320, SRX345 and SRX380 platforms. [PR1878769](#)
- SRX Series Firewall with IOC3 triggers a temperature alert on the FPC 2 PLX PCI Express Switch Chip. [PR1883027](#)
- Policy match failure for VXLAN EVPN type-5 cross vrf traffic. [PR1884150](#)
- Alarms for high usage in /var partition are not generated. [PR1886757](#)
- SRX4600 node 1 enters hardware failure during upgrade. [PR1887000](#)
- The XE interfaces of SRX380 platform with 1 G SFP (fiber) are flapping continuously when LACP is enabled. [PR1889549](#)
- IDP installation update fails on secondary node in SRX Series Firewall chassis cluster. [PR1890791](#)
- Security logging does not work on SRX Series Firewall when the stream host is configured using FQDN that resolves to an IPv6 address. [PR1892867](#)
- SRX Series Firewall configured with a native VLAN ID other than 1 experienced DHCP assignment issues and ARP resolution failures to the default gateway. [PR1893957](#)
- Packet drops are observed on SRX380 platforms in packet mode. [PR1897579](#)
- Skipping processing for smaller files affected fragmented packets. [PR1899463](#)
- Packet Forwarding Engine process might stop where PKI and SSL-Proxy services are configured. [PR1901098](#)
- SRX4700 performance improvement for TYPE5 VXLAN scenario. [PR1901678](#)
- SSL profiles not getting synchronize with Packet Forwarding Engine on SRX Series Firewall. [PR1903639](#)
- RIB and the FIB inconsistency results in traffic loss in IPsec scenario with st0 interface configured. [PR1908681](#)
- Link status of disabled SFP-T port becomes up after rebooting on SRX1500. [PR1910445](#)
- SRX Series Firewall Packet Forwarding Engine process might stop if nexthop limit is reached. [PR1911845](#)

Infrastructure

- SRX5600 and SRX5800 SPC3 are going offline after software upgrade. [PR1879079](#)

Intrusion Detection and Prevention (IDP)

- Not able to download IDP signature through routing instance. [PR1883645](#)

J-Web

- Software upgrade through J-Web does not work on specific SRX Series Firewall while using Upload Package option. [PR1916722](#)

Platform and Infrastructure

- In SRX Series Firewall high availability cluster, RGO failover to secondary node fails as srpxpf processes failed to reconnect to routing-engine on secondary. [PR1904267](#)

Routing Policy and Firewall Filters

- On SRX Series Firewall platform flowd process is generating core files. [PR1882193](#)
- Policy configuration failure causes the srpxpf process to stop responding. [PR1888456](#)
- Support IPv6 in address book address range. [PR1892746](#)
- SRX5K traffic disruption due to REPFE policy sync issues from FQDN and file-serialization Errors. [PR1894033](#)
- SRX Series Firewall and MX-SPC3: While Downgrading from 25.4/24.4R2/25.2R2/25.3R1 image with address book IPV4 range config, image installation validation is failing. [PR1899519](#)

Content Security

- Policy-name added to web filtering syslogs. [PR1879364](#)

VPNs

- Added more flags to "ike traceoptions flag" configuration statement. [PR1804885](#)
- The iked process might stop due to continuous activate and deactivate of VPN configuration. [PR1869875](#)
- IKE processing order causing gateway misconfiguration and tunnel failures. [PR1892539](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 170

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

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

For information about ISSU, see the [Chassis Cluster User Guide for Security Devices](#).

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

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

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

Junos OS Release Notes for vSRX

IN THIS SECTION

- [What's New | 172](#)
- [What's Changed | 176](#)
- [Known Limitations | 179](#)
- [Open Issues | 179](#)
- [Resolved Issues | 179](#)
- [Migration, Upgrade, and Downgrade Instructions | 180](#)

What's New

IN THIS SECTION

- [Application Identification \(AppID\) | 172](#)
- [Application Layer Gateways \(ALGs\) | 172](#)
- [Device Security | 173](#)
- [High Availability | 173](#)
- [Identity Aware Firewall | 174](#)
- [Junos OS API and Scripting | 174](#)
- [Junos Traffic Vision | 175](#)
- [Public Key Infrastructure \(PKI\) | 176](#)

Learn about new features introduced in this release for vSRX.

Application Identification (AppID)

- **Authentication support in proxy profiles (SRX Series, cSRX, and vSRX)**—You can enhance your network's security by configuring authentication within proxy profiles, allowing secure access to external feeds and services. By setting a username and password, you enable authenticated HTTPS communication through a proxy, supporting services such as SecIntel, AppID, Content Security, and PKID. This configuration prevents unverified data sources from accessing protected environments. To configure the feature, set a username and password in your proxy profile and use HTTPS to encrypt credentials, mitigating security vulnerabilities. Regularly update passwords and monitor for unauthorized access to maintain security integrity.

[See [Configuring SSL Proxy](#).]

Application Layer Gateways (ALGs)

- **Support for client identifier in forwarded DNS queries using experimental EDNS(0) option (NFX150, NFX250, NFX350, cSRX, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—You can enhance DNS query responses by utilizing the *Client ID in Forwarded DNS Queries* feature on SRX Series devices. This functionality supports precise client identification, essential for services such as parental control. You can configure SRX Series Firewalls to include client identifiers such as MAC addresses, IPv4, or IPv6 addresses in DNS queries. This feature, intended for controlled environments, generates targeted

DNS responses based on the originating device's identity, improving service accuracy and network efficiency.

[See [DNS ALG](#).]

Device Security

- **VRF-aware zone-based security policies with flow enhancements for Layer 3 VPN over EVPN-VXLAN and MPLS (SRX Series and vSRX 3.0)**—You can enforce security policies for each virtual routing and forwarding (VRF) instance by creating VRF-aware security zones. This approach helps improve scalability and intuitive policy management across EVPN-VXLAN and Layer 3 VPN (L3VPN) segments. You can define zones by VRF instance, not by VRF group, and implement intra-VRF or inter-VRF policies using the CLI or the management user interface.

[See [Security Policies with VRF-Aware Security Zones](#), [Flow Management in SRX Series Devices Using VRF Routing Instance](#), [show security flow session](#), and [security-zone](#).]

- **Group-based policy in VXLAN architecture (SRX Series and vSRX)**—Use a group-based policy (GBP) to create microsegmentation in VXLAN architecture by defining application-centric policies. Associate endpoints with tags that identify business functions to manage network access and direct traffic between endpoint groups. Enforce granular access control using new match options for source and destination tags. This approach strengthens security and simplifies policy enforcement across campus networks.

[See [Group-Based Policies in VXLAN Environments](#).]

- **FQDN ID for enhanced policy management and dynamic IP resolution (SRX Series, cSRX, and vSRX)**—Use unique fully qualified domain name (FQDN) ID mappings to manage frequent IP address changes. The system stores each FQDN's identifier (ID) in the Routing Engine and Packet Forwarding Engine. This FQDN ID storage enables quick lookups without constant policy updates, improving stability. This feature runs by default.

[See [DNS Snooping for Security Policies](#).]

High Availability

- **Dual-path ICL for MNHA in cloud environments (vSRX)**—You can configure dual-path interchassis links (ICLs) with aggregated Ethernet and loopback (lo0) interfaces for public and private cloud Multinode High Availability (MNHA). We recommend loopback interfaces in Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) because of aggregated Ethernet interface limitations. In private clouds with kernel-based virtual machine (KVM) or VMware EXSi, configure aggregated Ethernet interfaces for flexible traffic distribution. Use five-tuple hashing for optimal load balancing across Packet Forwarding Engines. This approach improves efficiency and reliability in MNHA.

[See [Multinode High Availability Support for vSRX Virtual Firewall Instances](#).]

- **Selective session synchronization for MNHA (SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—Use selective session synchronization for Multinode High Availability (MNHA) during cold and hot synchronization to optimize performance, reduce redundant state replication, and maintain fine-grained control over synchronization durations. Disable synchronization for short-lived traffic or defer synchronization with a minimum age. You can configure selective session synchronization using default or user-defined flow profiles. You can disable session synchronization for short-lived sessions by setting `session-sync disabled` or delay synchronization based on session age by adjusting `session-sync-min-age`.

[See [Selective Session Synchronization for Multinode High Availability](#).]

Identity Aware Firewall

- **Optimization of IPC messages for unified access control (UAC) authentication entries (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)**—Optimize system performance by using IPC message communication between the Routing Engine and Packet Forwarding Engine. Group multiple role entries into one message to improve efficiency. Use the `clear services unified-access-control authentication-table` command to refresh the UAC authentication table for accurate role and user data.

[See [clear services unified-access-control authentication-table](#).]

- **User identity through HTTP XFF header (SRX Series Firewalls, and vSRX3.0)**—You can identify users behind proxies by extracting the originating client IP from HTTP X-Forwarded-For (XFF) or Forwarded headers. Avoid relying on the packet source IP. Use this method to prevent misattribution from proxy addresses and to improve policy enforcement, logging, and analytics accuracy. Use the `set services user-identification forward-header-lookup` command to enable the XFF header feature.

[See [forward-header-lookup](#), [show services user-identification forward-header-statistics](#), and [clear services user-identification forward-header-statistics](#).]

Junos OS API and Scripting

- **Support for libslax 3.1.6 and SLAX version 1.3 (EX2300, EX2300-C, EX2300-MP, EX2300-VC, EX3400, EX3400-VC, EX4000-8P, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4100-24MP, 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, EX4100-H-12MP, EX4100-H-12MP-DC, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4300-MP, EX4300VC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4600-VC, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, cSRX, SRX300, SRX320, SRX340, SRX345,**

SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0—We've upgraded the libslax library to version 3.1.6, which corresponds to SLAX version 1.3. The upgraded library includes:

- Slax processor enhancements including a new mode, additional options, and simplified argument parsing
- New libslax extension library functions
- Improved SLAX syntax options
- New SLAX functions and enhancements to existing functions and statements
- Hexadecimal numbers support in SLAX scripts

This update aligns SLAX processing with contemporary scripting conventions, ensuring efficient, readable scripting while maintaining backward compatibility.

[See [libslax Distribution Overview](#).]

Junos Traffic Vision

- **Junos Traffic Vision (SRX Series Firewall and vSRX 3.0)**—We now support Junos Traffic Vision (previously known as Jflow) on SRX Series Firewall. Junos Traffic Vision uses security policies and logical systems to generate flow records that capture addresses, packet counts, and byte counts. NAT44, NAT64, and NAT66 session template records include new information using J-Flow version 9 and IPFIX.

The new elements are:

IANA IPFIX ID	Field Name	Size (Bytes)
152	flowStartMilliseconds	32
6	tcp_flags	16
10	input_snmp	32
14	output_snmp	32
95	application_id	32
234	ingressVRFID	32

(Continued)

IANA IPFIX ID	Field Name	Size (Bytes)
235	egressVRFID	32

Specify only the required template events and assign a predefined value (256 through 65535) for reuse.

Templates also support non-NAT firewall IPv4 and IPv6.

[See [Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250](#) and [Junos Traffic Vision Support on MS-MIC and MS-MPC](#).]

Public Key Infrastructure (PKI)

- **HTTPS support for PKI (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)**—With HTTPS support for PKI, you can enhance the security of certificate management operations. This feature establishes secure communication channels for SCEP enrollment and CRL revocation, protecting sensitive information. The PKI process dynamically selects HTTP or HTTPS based on configured URLs, providing flexibility and secure transmissions.

[See [PKI Components in Junos OS](#).]

What's Changed

IN THIS SECTION

- [Authentication and Access Control | 177](#)
- [Dynamic Host Configuration Protocol | 177](#)
- [Infrastructure | 177](#)
- [Network Management | 177](#)
- [Platform and Infrastructure | 177](#)
- [User Interface and Configuration | 178](#)
- [VPNs | 178](#)

Learn about what changed in this release for vSRX.

Authentication and Access Control

- For push-to-identity-management to successfully push the authentication entry to JIMS, you must configure JIMS and verify that JIMS status is online.

[See [push-to-identity-management](#).]

Dynamic Host Configuration Protocol

- The `show dhcp-security vlans detail` command now correctly displays the Interface Flag column, which shows the trust status for each interface. Flag values are hexadecimal bitmasks indicating: 0x1 (trusted), 0x2 (untrusted), 0x4 (no option82), 0x8 (no option37), 0x10 (no option18), 0x20 (no option16), 0x40 (no dhcpv6 options), and 0x80 (no option79).

[See [show dhcp-security vlans](#).]

Infrastructure

- You can now boot vSRX 3.0 with either UEFI or BIOS.

Network Management

- **IPv6 DNS resolution option in security log stream configuration (SRX Series Firewalls and vSRX3.0)**—You can enable the `prefer-ipv6-dns` option under the `show security log stream s1` host configuration hierarchy to prioritize IPv6 address resolution for DNS queries. This option ensures that IPv6 addresses are used instead of the default IPv4 addresses. This configuration enhances IPv6 network compatibility and supports environments that require IPv6 addressing.

Platform and Infrastructure

- **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](#).]

- You can now enable zeroization on a vSRX 3.0 Virtual Firewall using CLI to destroy Critical Security Parameters (CSPs). Run the `request system zeroize` command to zeroize the system configuration and

keys. When you run this command all the configuration information is removed, and the key values are reset and the vSRX 3.0 firewall is reverted to factory defaults after reboot.

User Interface and Configuration

- **Generate genstate YANG modules on Junos devices**—You can use `show system schema operational` command or equivalent RPC to generate the genstate YANG modules in the specified output directory on a device.

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

VPNs

- **Default installation of junos-ike package on additional platforms (SRX1500, SRX4100, SRX4200, SRX4600, and vSRX 3.0)**—The `junos-ike` package is installed by default on SRX1500, SRX4100, SRX4200, SRX4600, and vSRX3.0 firewalls, ensuring the default support for `iked` process for IPsec VPN service. This aligns with the existing default installation of the package on SRX5000 line with Routing Engine 3 (SRX5K-SPC3 with RE3). You can delete the `junos-ike` package using the command `request system software delete junos-ike`. This runs the `kmd` process on these firewalls, allowing flexible management of your security infrastructure.

[See [IPsec VPN Overview](#).]

- **Deprecation of weak algorithms in IPsec VPN (SRX Series and vSRX 3.0)**—We've deprecated the weak algorithms in IKE and IPsec proposals. You'll no longer be able to use the following algorithms:

Table 14: Deprecated Junos CLI Options

Type	Algorithm	Junos CLI Statement
Encryption Algorithm in IKE Proposal	des-cbc and 3des-cbc	<code>set security ike proposal <i>name</i> encryption-algorithm</code>
Authentication Algorithm in IKE Proposal	md5 and sha1	<code>set security ike proposal <i>name</i> authentication-algorithm</code>
DH Group in IKE Proposal	group1, group2, and group5	<code>set security ike proposal <i>name</i> dh-group</code>
Encryption Algorithm in IPsec Proposal	des-cbc and 3des-cbc	<code>set security ipsec proposal <i>name</i> encryption-algorithm</code>

Authentication Algorithm in IKE Proposal	hmac-md5-96 and hmac-sha1-96	set security ipsec proposal <i>name</i> authentication-algorithm
--	------------------------------	---

You will receive a warning message if you configure these deprecated algorithms explicitly. As an alternative, we recommend that you configure the stronger algorithms to enhance the security in IPsec VPN.

[See [proposal \(Security IKE\)](#), and [proposal \(Security IPsec\)](#).]

- **Support for hmac-sha-384/512 authentication in PMI (SRX Series Firewalls and vSRX 3.0)**—You can configure hmac-sha-384 and hmac-sha-512 authentication algorithms with PowerMode IPsec (PMI) when running IPsec VPN with the iked process.

[See [PowerMode IPsec](#).]

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

There are no known issues 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.

Resolved Issues

IN THIS SECTION

- [General Routing](#) | 180

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.

General Routing

- The flowd process might stop on all Junos SRX Series Firewall in multicast scenario with PIM. [PR1877771](#)
- Policy match failure for VXLAN EVPN type-5 cross VRF traffic. [PR1884150](#)
- On VSRX 3.0 support for zeroization command is enabled. [PR1886948](#)
- IDP installation update fails on secondary node in SRX Series Firewall chassis cluster. [PR1890791](#)
- Clear security log reports with time interval support. [PR1892154](#)
- SSL profiles not getting synchronize with Packet Forwarding Engine on SRX Series Firewall. [PR1903639](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases](#) | 186

This section contains information about how to upgrade Junos OS for vSRX using the CLI. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

You also can upgrade to Junos OS Release for vSRX using J-Web (see [J-Web](#)) or the Junos Space Network Management Platform (see [Junos Space](#)).

Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Releases 17.4, 18.1, 18.2, 18.3, 18.4, 19.1, 19.2 and 19.4 is supported.

The following limitations apply:

- Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Release 19.3 and higher is not supported. For upgrade between other combinations of Junos OS Releases in vSRX and vSRX 3.0, the general Junos OS upgrade policy applies.
- The file system mounted on /var usage must be below 14% of capacity.

Check this using the following command:

```
show system storage | match " /var$" /dev/vtbd1s1f
2.7G      82M      2.4G      3% /var
```

Using the request system storage cleanup command might help reach that percentage.

- The Junos OS upgrade image must be placed in the directory /var/host-mnt/var/tmp/. Use the request system software add /var/host-mnt/var/tmp/<upgrade_image>
- We recommend that you deploy a new vSRX virtual machine (VM) instead of performing a Junos OS upgrade. That also gives you the option to move from vSRX to the newer and more recommended vSRX 3.0.
- Ensure to back up valuable items such as configurations, license-keys, certificates, and other files that you would like to keep.



NOTE: For ESXi deployments, the firmware upgrade from Junos OS Release 15.1X49-Dxx to Junos OS releases 17.x, 18.x, or 19.x is not recommended if there are more than three network adapters on the 15.1X49-Dxx vSRX instance. If there are more than three network adapters and you want to upgrade, then we recommend that you either delete all the additional network adapters and add the network adapters after the upgrade or deploy a new vSRX instance on the targeted OS version.

Upgrading Software Packages

To upgrade the software using the CLI:

1. Download the **Junos OS Release 25.4R1 for vSRX .tgz** file from the [Juniper Networks website](#). Note the size of the software image.
2. Verify that you have enough free disk space on the vSRX instance to upload the new software image.

```
root@vsrx> show system storage
```

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/vtbd0s1a	694M	433M	206M	68%	/
devfs	1.0K	1.0K	0B	100%	/dev

/dev/md0	1.3G	1.3G	0B	100%	/junos
/cf	694M	433M	206M	68%	/junos/cf
devfs	1.0K	1.0K	0B	100%	/junos/dev/
procfs	4.0K	4.0K	0B	100%	/proc
/dev/vtbd1s1e	302M	22K	278M	0%	/config
/dev/vtbd1s1f	2.7G	69M	2.4G	3%	/var
/dev/vtbd3s2	91M	782K	91M	1%	/var/host
/dev/md1	302M	1.9M	276M	1%	/mfs
/var/jail	2.7G	69M	2.4G	3%	/jail/var
/var/jails/rest-api	2.7G	69M	2.4G	3%	/web-api/var
/var/log	2.7G	69M	2.4G	3%	/jail/var/log
devfs	1.0K	1.0K	0B	100%	/jail/dev
192.168.1.1:/var/tmp/corefiles	4.5G	125M	4.1G	3%	/var/crash/ corefiles
192.168.1.1:/var/volatile	1.9G	4.0K	1.9G	0%	/var/log/host
192.168.1.1:/var/log	4.5G	125M	4.1G	3%	/var/log/hostlogs
192.168.1.1:/var/traffic-log	4.5G	125M	4.1G	3%	/var/traffic-log
192.168.1.1:/var/local	4.5G	125M	4.1G	3%	/var/db/host
192.168.1.1:/var/db/aamwd	4.5G	125M	4.1G	3%	/var/db/aamwd
192.168.1.1:/var/db/secinteld	4.5G	125M	4.1G	3%	/var/db/secinteld

3. Optionally, free up more disk space, if needed, to upload the image.

```

root@vsrx> request system storage cleanup
List of files to delete:
Size Date      Name
11B Sep 25 14:15 /var/jail/tmp/alarmd.ts
259.7K Sep 25 14:11 /var/log/hostlogs/vjunos0.log.1.gz
494B Sep 25 14:15 /var/log/interactive-commands.0.gz
20.4K Sep 25 14:15 /var/log/messages.0.gz
27B Sep 25 14:15 /var/log/wtmp.0.gz
27B Sep 25 14:14 /var/log/wtmp.1.gz
3027B Sep 25 14:13 /var/tmp/BSD.var.dist
0B Sep 25 14:14 /var/tmp/LOCK_FILE
666B Sep 25 14:14 /var/tmp/appidd_trace_debug
0B Sep 25 14:14 /var/tmp/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 25.4R1 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-vsr-x86-64-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE.tgz /var/crash/corefiles/
```

5. From operational mode, install the software upgrade package.

```
root@vsrx> request system software add /var/crash/corefiles/junos-vsr-x
x86-64-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE.tgz no-copy no-validate reboot
Verified junos-vsr-x86-64-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING: This package will load JUNOS 20.4 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-vsr-x-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE.tgz
Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31 junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31 junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz
Setting up Junos host applications for installation ...
```

```

=====
Host OS upgrade is FORCED
Current Host OS version: 3.0.4
New Host OS version: 3.0.4
Min host OS version required for applications: 0.2.4
=====

Installing Host OS ...
upgrade_platform: -----
upgrade_platform: Parameters passed:
upgrade_platform: silent=0
upgrade_platform: package=/var/tmp/junos-srx-mr-vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-
linux.tgz
upgrade_platform: clean install=0
upgrade_platform: clean upgrade=0
upgrade_platform: Need reboot after staging=0
upgrade_platform: -----
upgrade_platform:
upgrade_platform: Checking input /var/tmp/junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz ...
upgrade_platform: Input package /var/tmp/junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz is valid.
upgrade_platform: Backing up boot assets..
cp: omitting directory '.'
bzImage-intel-x86-64.bin: OK
initramfs.cpio.gz: OK
version.txt: OK
initrd.cpio.gz: OK
upgrade_platform: Checksum verified and OK...
/boot
upgrade_platform: Backup completed
upgrade_platform: Staging the upgrade package - /var/tmp/junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz..
./
./bzImage-intel-x86-64.bin
./initramfs.cpio.gz
./upgrade_platform
./HOST_COMPAT_VERSION
./version.txt
./initrd.cpio.gz
./linux.checksum
./host-version
bzImage-intel-x86-64.bin: OK
initramfs.cpio.gz: OK

```

```

version.txt: OK
upgrade_platform: Checksum verified and OK...
upgrade_platform: Staging of /var/tmp/junos-srx-mr-
vsrx-20.4-2020-10-12.0_RELEASE_20.4_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 25.4R1 for vSRX.



NOTE: Starting in Junos OS Release 17.4R1, upon completion of the vSRX image upgrade, the original image is removed by default as part of the upgrade process.

6. Log in and use the show version command to verify the upgrade.

```

--- JUNOS 20.4-2020-10-12.0_RELEASE_20.4_THROTTLE Kernel 64-bit
JNPR-11.0-20171012.170745_fbsd-
At least one package installed on this device has limited support.
Run 'file show /etc/notices/unsupported.txt' for details.
root@:~ # cli
root> show version
Model: vsrx
Junos: 20.4-2020-10-12.0_RELEASE_20.4_THROTTLE
JUNOS OS Kernel 64-bit [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs [20171012.170745_fbsd-builder_stable_11]
JUNOS OS runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS OS time zone information [20171012.170745_fbsd-builder_stable_11]

```

```

JUNOS OS libs compat32 [20171012.170745_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20171012.170745_fbsd-builder_stable_11]
JUNOS py extensions [20171017.110007_ssd-builder_release_174_throttle]
JUNOS py base [20171017.110007_ssd-builder_release_174_throttle]
JUNOS OS vmguest [20171012.170745_fbsd-builder_stable_11]
JUNOS OS crypto [20171012.170745_fbsd-builder_stable_11]
JUNOS network stack and utilities [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Web Management Platform Package [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS common platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS mtx network modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx Data Plane Crypto Support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Online Documentation [20171017.110007_ssd-builder_release_174_throttle]
JUNOS jail runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS FIPS mode utilities [20171017.110007_ssd-builder_release_174_throttle]

```

Validating the OVA Image

If you have downloaded a vSRX .ova image and need to validate it, see [Validating the vSRX .ova File for VMware](#).

Note that only .ova (VMware platform) vSRX images can be validated. The .qcow2 vSRX images for use with KVM cannot be validated the same way. File checksums for all software images are, however, available on the download page.

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

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

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

- Extended End of Life (EEOL) releases have engineering support for 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 15: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/Downgrade to subsequent 3 releases	Upgrade/Downgrade to subsequent 2 EEOL releases
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 | 189
- Creating a Service Request with JTAC | 190

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <https://www.juniper.net/content/dam/www/assets/resource-guides/us/en/jtac-user-guide.pdf>.
- Product warranties—For product warranty information, visit <https://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.

- Visit [Juniper Support Portal: Case Management, Product Support & More](#)
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

Revision History

24 September 2026—Revision 18, Junos OS Release 25.4R1.

04 September 2026—Revision 17, Junos OS Release 25.4R1.

19 August 2026—Revision 16, Junos OS Release 25.4R1.

13 August 2026—Revision 15, Junos OS Release 25.4R1.

17 July 2026—Revision 14, Junos OS Release 25.4R1.

13 July 2026—Revision 13, Junos OS Release 25.4R1.

25 June 2026—Revision 12, Junos OS Release 25.4R1.

4 June 2026—Revision 11, Junos OS Release 25.4R1.

7 May 2026—Revision 10, Junos OS Release 25.4R1.

4 May 2026—Revision 9, Junos OS Release 25.4R1.

30 April 2026—Revision 8, Junos OS Release 25.4R1.

28 April 2026—Revision 7, Junos OS Release 25.4R1.

13 March 2026—Revision 6, Junos OS Release 25.4R1.

23 February 2026—Revision 5, Junos OS Release 25.4R1.

4 February 2026—Revision 4, Junos OS Release 25.4R1.

30 January 2026—Revision 3, Junos OS Release 25.4R1.

22 January 2025—Revision 2, Junos OS Release 25.4R1.

22 December 2025—Revision 1, Junos OS Release 25.4R1.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2026 Juniper Networks, Inc. All rights reserved.