Release Notes

Published
2026-08-31

Junos OS Release 26.2R1

Juniper Networks | Engineering Simplicity

Table of Contents

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 26.2R1. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can find release notes for all Junos OS releases at Junos OS Documentation | Juniper Networks.

Junos OS Release Notes for ACX Series

IN THIS SECTION

What's New

IN THIS SECTION

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

Junos Telemetry

OpenConfig YANG model support for BGP (ACX5448, MX204, MX240, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vRR)—Junos telemetry now supports the OpenConfig YANG model for BGP, providing a vendor-neutral framework for configuration and monitoring BGP. Retrieve BGP-specific operational and state data by subscribing to sensors under resource path hierarchies such as `inline communities`, `community-set`, `match-community-set`, `match-ext-community-set`, `peer-group`, `neighbor`, `policy-definition`, `administrative distance`, and `as-path-set`. OpenConfig-defined default behavior for BGP policies is implemented, with `community-level` policy options supporting `any-match` and `invert-match` configurations.

Use the Junos YANG Data Model Explorer to view supported sensor paths and platforms. The supported OpenConfig version for the openconfig-network-instance model is 4.4.1.

The following table lists the supported YANG models and their versions:

BGP YANG Model	Version
openconfig-routing-policy.yang	3.4.2
openconfig-bgp-common-multiprotocol.yang	9.9.1
openconfig-bgp-common-structure.yang	9.9.1
openconfig-bgp-common.yang	9.9.1
openconfig-bgp-errors.yang	6.1.0
openconfig-bgp-global.yang	9.9.1
openconfig-bgp-neighbor.yang	9.9.1
openconfig-bgp-peer-group.yang	9.9.1
openconfig-bgp-policy.yang	8.1.0
openconfig-bgp-types.yang	6.1.0

(Continued)

BGP YANG Model	Version
openconfig-bgp.yang	9.9.1
openconfig-types.yang	1.0.0

[See Junos YANG Data Model Explorer.]

Enhancement of BGP XPath policy (ACX5448, MX480, MX960, MX2020, MX10004, and vRR)—Add OpenConfig BGP XPath enhancements for default policy and policy definitions.

[See Junos YANG Data Model Explorer.]

User Access and Authentication

Enhancement: TLS Support for TACACS+ (TACACS+/TCP Secure Transport) (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—This feature enhances the existing TACACS+ over TCP (TACACS+/TCP) implementation by adding support for Transport Layer Security (TLS), enabling secure TACACS+ communication. With this update, the device-side TACACS+ client is extended to establish TLS sessions with TACACS+ servers, ensuring encryption of authentication, authorization, and accounting (AAA) traffic. The enhancement includes TLS handshake handling, server certificate validation, and secure session establishment, while maintaining backward compatibility with non-TLS configurations. This improves overall security by protecting sensitive data from interception and aligns TACACS+ communication with modern security standards.

Additional Features

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

Support for Quantum Buffer in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP,

EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptographic agility.

[See The Quantum Buffer and moduli.]

Support for Shor-resistant and other default key exchange algorithms in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSH supports the Shor's algorithm-resistant hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm. And, by default, SSH supports the `dh-group16-sha512` and `dh-group18-sha512` Diffie-Hellman (DH) group key exchange algorithms.

[See key-exchange.]

What's Changed

IN THIS SECTION

Learn about what's changed 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 OS Problem Report Search application.

EVPN

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit protocols evpn]` hierarchy level only on platforms that support a default switch instance. To enable

these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the [edit routing-instances name protocols evpn] hierarchy level.

[See Preference-Based DF Election.]

- **XML output tags changed for show l2-learning evpn mcnh-info command**—We've changed the XML output for this command. The XML <l2ald-mcnh-entry-list-headers> element level under the <l2ald-mcnh-entry> element has been removed. The child output tags that were under the <l2ald-mcnh-entry-list-headers> element are now direct child tags of the <l2ald-mcnh-entry-list> element.

General Routing

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the [edit services port-extender satellite <name> device-model] hierarchy, you can use the '?' to see a list of supported platforms. We have also added the skip-lo0-config statement under [edit service port-extender jnu] hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added cp-mtu under the [edit service port-extender satellite <name>] hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

  [See port-extender (Services).]

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for auto-configuration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address ff02::1. The RA source is the router's link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the show route extensive expanded-nh command.

  [See show route extensive <route> expanded-nh.]

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)               - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)        -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- **Command to view interface-set information (ACX Series, MX Series)**—Use the `show subscribers interfaces interface-set` command to view interface-set information.

  [See show subscribers interfaces.]

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **Per-fan entries for fan trays in SNMP MIB tables (PTX10008)**—You can monitor individual fans in a fan tray through distinct entries in SNMP MIB tables. This enhancement addresses the earlier limitation that SNMP MIB does not list individual fans.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.

  - A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Member Link Limit on Adaptive Aggregated Ethernet (PTX Series)**—A maximum of 64 member links are configurable on an aggregated Ethernet (AE) bundle with adaptive load balancing. Configuring more than 64 member links will result in a commit failure.

  [See Understanding Aggregated Ethernet Load Balancing.]

- **ARP DDoS protection configuration restriction for unicast and broadcast policers (ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509)**—For ARP DDoS protection on devices running Junos OS Evolved in the ACX family, only the aggregate policer option is supported and effective under the system ddos-protection protocols arp hierarchy. The ucast and bcast ARP policer options are not supported on this Platform and are now treated as no-operations: if you configure these options, the system ignores the corresponding configuration blocks and reports a warning message "Warning: configuration block ignored: unsupported platform" indicating that the configuration block is ignored for the platform.

- Write cache disabled alarms no longer raised—The "Disk 2 Write Cache disabled" alarms were reported in the chassis alarms due to a missing SSD model in Junos OS supported device table. The support is now added, and alarms are no longer reported.

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading, copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

[See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:

  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.

  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.

  - The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

  These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

  [See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Network Management and Monitoring

- **Improved ZTP SNMP configuration handling (Junos)**—We have reordered the SNMP configuration commands to optimize the Zero Touch Provisioning (ZTP) process. The `set snmp engine-id local` statement now processes before `set snmp v3 usm local-engine user`, ensuring SNMPv3 USM authentication and privacy keys generation using the configured local engine ID instead of an auto-derived default. This prevents SNMPv3 query authentication failures from engine-ID resets.

  [See Configure Engine ID on SNMPv3 and Create SNMPv3 Users.]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

[See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **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.]

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **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.

User Interface and Configuration

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:

  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`

- Changed the RPC to rename the request tag from <get-genstate-xpath-cli-command-mapping> to <get-genstate-path-cli-command-mapping> and rename the <xpath> option to <genstate-path>
- Changed the XML output to rename the root element from <genstate-xpath-cli-command-mapping-information> to <genstate-path-cli-command-mapping-information> and rename the <xpath> element to <genstate-path>

[See show system data-models.]

- **Introduction of remove operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents <userinput>statement not found</userinput> warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.
- **CLI options and RPC formats to retrieve `genstate` information and XPath expressions for operational commands**—A gNMI telemetry collector can subscribe to `genstate` resource paths to query for the corresponding state data on a device. We've added CLI pipe filter functions and RPC formats that enable you to discover available `genstate` resource paths and the associated `genstate` YANG module for a given show command. You can also retrieve the XPath expression for any node in command output. The options include:
  - `| display genstate subscription-paths`—Display all resource paths available for subscription in the `genstate` YANG module associated with the given command.
  - `| display genstate yang-module-name`—Display the name of the `genstate` YANG module that includes the data model and resource paths for the given command.
  - `| display xml genstate-path`—Display the command output in XML and include the `junos:genstate-path` attribute for every emitted tag that is available for subscription. The attribute value is the `genstate` subscription path. For RPCs, use the `format='xml-genstate-path'` attribute in the request tag.
  - `| display xml rpc response-paths`—List all possible response paths under the equivalent RPC's top-level response tag.
  - `| display xml xpath`—Display the command output in XML and include the `junos:data-xpath` attribute for each tag in the output. The attribute value is the XPath expression that selects that specific node. For RPCs, use the `format='xml-xpath'` attribute in the request tag.

[See | (pipe) and Junos Genstate YANG Data Models.]

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

Known Limitations

IN THIS SECTION

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

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

General Routing

- JUNOS_REG: ACX5448-M: While verifying traffic convergence with FRR on DUT traffic loss duration is more than expected

  PR1927951

Open Issues

IN THIS SECTION

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

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

General Routing

- For ACX5448, MX204 and MX2008 "VM Host-based" platforms, starting with Junos 21.4R1 or later, ssh and root login is required for copying line card image (chspmb.elf for MX2008) from Junos VM to Linux host during installation. The ssh and root login are required during installation. Use "deny-password" instead of "deny" as default root-login option under ssh config to allow internal trusted communication. Ref https://kb.juniper.net/TSB18224 PR1629943
- On ACX710 platforms running Junos OS, the acx-arm-feb process reaches 100 percent CPU utilization after a system upgrade or reboot when Precision Time Protocol (PTP) or Synchronous Ethernet (SyncE) is configured and the associated interfaces are enabled. This condition results in control plane degradation and traffic impact.

  PR1900889

Resolved Issues

IN THIS SECTION

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

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

General Routing

- On all ACX5448 platform, during a CCC(circuit cross-connect)-to-MPLS primary and protect path next-hop swap event, the in-labels are also being swapped. However, these labels were already swapped during earlier MPLS-to-CCC events. This additional, unintended swap by the PFE leads to traffic disruption. PR1911023

- When MTU set as 1500, transient packets were redirected to the host path instead of being forwarded directly. Thus causing ping latency.This issue was observed after the changes introduced in PR1775703 and the changes are committed in 24.1. All the releases after 24.1 will face the issue, previous releases will not have any issue. PR1903922
- On all Junos platforms, if an interface is configured with "input-vlan-map tag-protocol-id" and the associated "ethernet-switch-profile tag-protocol-id" is deleted while the VLAN mapping is still active, the system may incorrectly allow the configuration commit. If the device is subsequently rebooted or the chassis is restarted, the device may fail to recover and remain down, resulting in traffic impact. PR1933501
- On ACX5448 platform, a high amount of disk I/O leads to a deadlock, which triggers a panic reboot and a vmcore file is generated. During device restart, all services are temporarily unavailable. PR1948145

Layer 2 Features

- Traffic drop seen in VPLS stream post "clear ospf neighbor" command. PR1911208

Routing Protocols

- NGMVPN multicast traffic is dropped post changing the route distinguisher and route target on L3VPN instances. This issue is intermittent and occurs when RD and RT are changed in all routing instances while traffic is active. In some cases, the issue is observed only in specific instances. When the issue occurs, traffic drops for the inclusive tunnel, affecting all flows on the inclusive tunnel. Recovery may not occur unless external events, such as a route change, take place. PR1909265

Junos OS Release Notes for cRPD

IN THIS SECTION

What's New

IN THIS SECTION

Learn about new features introduced in this release for cRPD.

Junos Telemetry

Periodic telemetry for interface, process, and memory sensors (OpenConfig and native) (cRPD)—You can stream interface, process, and memory telemetry from CRPD using OpenConfig-compliant sensors through XMLproxyd. Use GNMI periodic subscriptions for interface state, counters, and error and drop XPaths. Include process XPaths (PID, name, CPU utilization) and native memory statistics (physical, active, free). Filter by interface paths to focus data. Verify and monitor subscriptions with <codeph>show agent sensors</codeph> and <codeph>show network-agent statistics detail</codeph>. To reduce CPU usage in the cRPD container when streaming these sensors in periodic mode, choose minute-level intervals and selective sensors. CRPD supports only data obtainable from the host system.

[See Junos YANG Data Model Explorer.]

Additional Features

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

Support for firewall filter and port mirroring (cRPD)—Use firewall filtering and port mirroring to protect routers from external incidents. Define and apply interface filtering policies to restrict unwanted traffic, and create mirroring sessions to forward selected packets to a monitoring device for inspection.

[See Configuring Port Mirroring and Firewall Filters Overview.]

Support for ping (cRPD)—Use the `ping` command from the configuration mode to troubleshoot the network traffic. Run `ping` with a destination address to send ICMP echo requests, verify connectivity, and assess latency and loss.

[See Debug cRPD Application.]

What's Changed

IN THIS SECTION

Learn about what's changed in this release for cRPD.

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

EVPN

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit protocols evpn]` hierarchy level only on platforms that support a default switch instance. To enable

these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the [edit routing-instances name protocols evpn] hierarchy level.

[See Preference-Based DF Election.]

- **XML output tags changed for show l2-learning evpn mcnh-info command**—We've changed the XML output for this command. The XML <l2ald-mcnh-entry-list-headers> element level under the <l2ald-mcnh-entry> element has been removed. The child output tags that were under the <l2ald-mcnh-entry-list-headers> element are now direct child tags of the <l2ald-mcnh-entry-list> element.

General Routing

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the [edit services port-extender satellite <name> device-model] hierarchy, you can use the '?' to see a list of supported platforms. We have also added the skip-lo0-config statement under [edit service port-extender jnu] hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added cp-mtu under the [edit service port-extender satellite <name>] hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

  [See port-extender (Services).]

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for auto-configuration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address ff02::1. The RA source is the router's link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the show route extensive expanded-nh command.

  [See show route extensive <route> expanded-nh.]

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)               - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)        -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **Per-fan entries for fan trays in SNMP MIB tables (PTX10008)**—You can monitor individual fans in a fan tray through distinct entries in SNMP MIB tables. This enhancement addresses the earlier limitation that SNMP MIB does not list individual fans.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.

- A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Member Link Limit on Adaptive Aggregated Ethernet (PTX Series)**—A maximum of 64 member links are configurable on an aggregated Ethernet (AE) bundle with adaptive load balancing. Configuring more than 64 member links will result in a commit failure.

  [See Understanding Aggregated Ethernet Load Balancing.]

- Write cache disabled alarms no longer raised—The "Disk 2 Write Cache disabled" alarms were reported in the chassis alarms due to a missing SSD model in Junos OS supported device table. The support is now added, and alarms are no longer reported.

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading, copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

  [See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:

  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.

  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.

  - The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

[See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Network Management and Monitoring

- **Improved ZTP SNMP configuration handling (Junos)**—We have reordered the SNMP configuration commands to optimize the Zero Touch Provisioning (ZTP) process. The `set snmp engine-id local` statement now processes before `set snmp v3 usm local-engine user`, ensuring SNMPv3 USM authentication and privacy keys generation using the configured local engine ID instead of an auto-derived default. This prevents SNMPv3 query authentication failures from engine-ID resets.

  [See Configure Engine ID on SNMPv3 and Create SNMPv3 Users.]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **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.]

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **System Stability Improvement**—When BGP fails to bind to its listening port, the routing protocol process (RPD) now stays running and logs a warning instead of terminating after the maximum retry count, preventing a service outage. The `show bgp diagnostics overview` output includes a new `BGP listening port status of v4 and v6` field that reports `Active` when BGP is successfully listening on the configured port and `Failed` when the port is unavailable or occupied by another process. Error logging is also enhanced with the `BGP_LISTEN_STATUS` message: `BGP listen failed on port [X] after 64 attempts with error-%d for address family identifier (IPv4 or IPv6). Administrative clear needed.` This message is generated after `BGP_MAX_LISTEN_START_ATTEMPTS` failed attempts to bind to the listening port.

User Interface and Configuration

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:

  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`

  - Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`

  - Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

  [See show system data-models.]

- **Introduction of remove operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents <userinput>statement not found</userinput> warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

Known Limitations

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

Open Issues

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

Resolved Issues

There are no resolved issues in this release for cRPD.

Junos OS Release Notes for cSRX

IN THIS SECTION

What's New

IN THIS SECTION

Learn about new features introduced in this release for cSRX.

Application Security

HTTP/2 AppQoS based on DSCP with first-stream classification for consistent policy enforcement (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Enforce granular application QoS for HTTP/2 traffic using differentiated services code point (DSCP) based on first-stream classifications. The device applies the application classification of the first-stream to the entire HTTP/2 parent session and all subsequent streams. Packets in the session inherit a uniform AppQoS behavior derived from that initial classification. Using first-stream classification ensures that traffic is classified and prioritized by application type when multiple streams use different applications.

This approach integrates with traditional and unified policy frameworks, simplifying deployment without requiring additional configuration.

[See Application QoS.]

HTTP/2 inspection for IDP (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4120, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Intrusion detection and prevention (IDP) now inspects HTTP/2 traffic to enhance security on SRX Series Firewalls. IDP identifies HTTP/2 application traffic and applies relevant attack signatures during traffic inspection. This capability increases visibility into modern web application traffic and improves threat detection over HTTP/2 connections.

IDP inspection secures HTTP/2 traffic within connection sessions and stream sessions. It detects attacks and anomalies at the protocol level.

You can create custom signatures using HTTP/2 contexts such as frame length, frame type, flags, and stream ID.

[See Support for HTTP/2 Traffic Inspection.]

HTTP/2 inspection for SSL proxy (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSL proxy adds Application-Layer Protocol Negotiation (ALPN)-based HTTP/2 inspection with policy controls for TLS 1.2 and TLS 1.3. During the TLS handshake, the proxy negotiates the HTTP/2 identifier h2 with compliant peers and falls back to HTTP/1.1 when HTTP/2 is disabled or not allowed.

HTTP/2 is disabled by default. You can enable or disable it globally by using the *set services ssl proxy global-config http2 on|off* command. *D*isable HTTP/2 selectively per policy using the *set security policies from-zone <from-zone> to-zone <to-zone> policy <policy-name> then permit application-services ssl-proxy disable-http2* command.

The SSL proxy enforces protocols by stripping HTTP/3 ALPN, blocks prohibited TLS 1.2 cipher suites, rejects renegotiation, and terminates TLS 1.3 post handshake authentication.

[See SSL Proxy HTTP/2 Inspection with ALPN (TLS 1.2 and TLS 1.3).]

HTTP/2 protocol support (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—HTTP/2 protocol is a modern version of the HTTP protocol designed to improve application performance and efficiency by enabling multiple transactions over a single connection. Each connection logically divides into multiple streams, enabling your device to inspect, permit, or deny transactions independently while maintaining a single connection.

You can safely enable applications running over HTTP/2 without any additional configuration on your device.

[See HTTP/2 Protocol.]

Support for cleartext HTTP/2 inspection (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Inspect unencrypted HTTP/2 traffic with cleartext HTTP/2 inspection. The firewall identifies HTTP/2 connections early, manages them as single TCP connections, and tracks each HTTP/2 stream independently. The system performs Layer 7 inspection and enforces policies on each stream, providing granular visibility and control over individual transactions within the same connection.

To configure cleartext inspection for HTTP/2, use the CLI commands: `set security application-services http2 plain-text` and `set services ssl proxy global config http2 on`. Verify your configuration with `show security application-services status` and `show security application-services monitoring`.

[See HTTP/2 Inspection: Support for Cleartext HTTP/2 Traffic.]

Support for HTTP/2 traffic in Juniper ATP Cloud (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)—Juniper ATP Cloud supports HTTP/2 traffic, giving you clear, stream-level visibility into HTTP/2 communications. It helps you detect threats accurately, inspect traffic deeply, and apply security policies consistently, without slowing down network performance. This support works seamlessly with Advanced Anti-Malware (AAMW), antivirus, Security Intelligence (SecIntel), and Security Metadata Streaming (SMS).

[See show services advanced-anti-malware statistics, show services anti-virus statistics, and show services security-metadata-streaming.]

Support for HTTP/2 in web filtering and content filtering (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—The HTTP/2 protocol supports Content Security features such as web filtering and content filtering. These features bypass the HTTP/2 parent session and process child sessions directly. The system applies actions based on the existing configuration for specific sessions, ensuring accurate filtering and enhanced security.

[See Content Security Overview and HTTP/2 Protocol.]

Content Security

Deprecation of performance-mode and disable-dynapp-profile-selection statements (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We've deprecated the `performance-mode` and `disable-dynapp-profile-selection` statements from the [`edit security utm default-configuration web-filtering`] hierarchy level.

[See web-filtering.]

Support for dynamic file type signatures in content filtering (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use dynamic file type signatures to efficiently manage file type signature data. These signatures are stored in a structured JSON format on a cloud-based server, allowing you to download, install, and update them without upgrading the existing Junos OS on your device. Before this release, you had to wait for the next Junos OS upgrade to update the file type signature data. Use CLI commands to easily manage tasks such as downloading, installing, and rolling back file type signatures.

[See request security utm content-filtering file-signatures download, request security utm content-filtering file-signatures download status, request security utm content-filtering file-signatures install, request security utm content-filtering file-signatures install status, request security utm content-filtering file-signatures local-update, request security utm content-filtering file-signatures rollback, request security utm content-filtering file-signatures rollback status, request security utm content-filtering file-

signatures uninstall, request security utm content-filtering file-signatures uninstall status, and show security utm content-filtering file-signatures status.]

Flow-Based and Packet-Based Forwarding

Optimization of TCP Proxy (cSRX, SRX1600, SRX2300, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We can use SRX Series Firewall as a TCP Proxy. This configuration ensures consistent proxying and prevents implicit services from disrupting sessions. It enhances session reliability and security. To configure security policy with default (global TCP Profile), use the `set security policies from-zone trust to-zone untrust policy p1 then permit tcp-proxy` command.

To configure the TCP proxy profile, use the `set security flow tcp-proxy tcp-profile` command. This command helps in configuring the required parameters.

To verify the configuration, use the`show security flow tcp-proxy profiles` command.

[See TCP Sessions.]

Junos Telemetry

Support for genstate YANG data models (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Junos telemetry supports the genstate YANG data model, a YANG-based model autogenerated from operational state data. Genstate models expose a subset of `show` command outputs through gNMI subscriptions. This feature allows a gNMI telemetry collector to subscribe directly to specific resource paths in the genstate model to retrieve the required state data. The genstate models mirror selected hierarchies and fields from Junos `show` commands, allowing you to query information such as interface attributes, protocol status, system metrics, and more.

For example, you can retrieve the `evpn-instance` information by configuring the gNMI collector to subscribe to the following resource path:

```
/genstate/evpn-instance-information/evpn-instance
```

In this path, `evpn-instance` is a field displayed in the output of the `show evpn instance` command.

Genstate resource paths are prefixed with `/genstate/`. The genstate YANG files are published on Github. See Genstate YANG Data Models to view the supported show commands and the root paths.

[See Model-Driven Telemetry, Data Models, and Explore Sensor Paths.]

Junos XML API and Scripting

Event policy archive sites support specifying a routing instance (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When you configure an archive site for an event policy destination, you can specify the routing instance that the device uses to connect to the site. To configure the routing instance, include the `routing-instance` option at the [`edit event-options destinations` *`destination-name`* `archive-sites` *`url`*] hierarchy level.

[See Event Policy File Archiving and archive-sites (Event Policy).]

Network Address Translation (NAT)

Deterministic NAT with interim logging (cSRX, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—You can now configure syslog intervals for deterministic NAPT. The configurable syslog interval reduces logging requirements by ensuring consistent IP address and port mappings.

Use the `set security nat source pool <pool-name> port deterministic det-nat-configuration-log-interval` command to define intervals from 1800 through 86400 seconds.

[See Deterministic NAT and det-nat-configuration-log-interval.]

Dynamic scaling for persistent NAT capacity (cSRX, SRX1600, SRX2300, SRX4120, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—You can enhance persistent NAT capacity by dynamically scaling beyond the fixed, code-defined limit. Dynamic scaling in persistent NAT enables growth by up to a configurable number of additional entries, depending on the configuration and available runtime resources, in addition to the minimum baseline static bindings allowed on the device.

You must configure the `set security nat source persistent-nat increased-scale-limit <value>` command to increase the persistent NAT capacity.

[See Persistent NAT and NAT64, show security nat source persistent-nat-table, and increased-scale-limit.]

User Access and Authentication

Sequential DNS fallback for outbound SSH (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When a configured hostname resolves to multiple IP addresses via DNS, Junos OS automatically tries each address in sequence, regardless of IP family, until a connection succeeds or all addresses have been attempted. The `outbound-ssh` retry mechanism activates after all resolved addresses have been tried.

[See outbound-ssh.]

VPNs

Configurable PRF in IKEv2 with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—You can configure Internet Key Exchange version 2 (IKEv2) proposals with a specific pseudorandom function (PRF) algorithm independent of the authentication algorithm. Use the `prf-algorithm` option to configure the following algorithms at the [`edit security ike proposal` *`proposal-name`*] hierarchy level in the IPsec VPN service with the iked process:

- `prf-hmac-sha-256`
- `prf-hmac-sha-384`
- `prf-hmac-sha-512`
- `prf-hmac-sha1`

You cannot combine a PRF with Authenticated Encryption with Associated Data (AEAD) cipher suites such as AES-GCM or ChaCha20-Poly1305 algorithm.

[See prf-algorithm and proposal (Security IKE).]

Deprecation of certificates as authentication method in IKE (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We've deprecated support for the `certificates` option at the [`edit security ike proposal` *`proposal-name`* `authentication-method`] hierarchy level. However, you can continue to use the other authentication methods at this hierarchy level.

[See proposal (Security IKE).]

Multiple local certificates in IKE with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—You can specify multiple local certificates in an IKE policy on a device that runs the iked process for the IPsec VPN service. During IKE negotiation, the device selects the local certificate whose certificate chain matches the peer's certificate request. The device evaluates the certificates in their listed order. Use `set security ike policy` *`policy-name`* `certificate local-certificate [`*`cert1 cert2`*`]` to configure multiple local certificates.

[See certificate and IKE Policy with Local Certificate.]

RSA-PSS signature algorithm in IKEv2 with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Specify the RSA signature scheme for IKEv2 authentication with an RSA local certificate as per RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2*. Configure the following algorithms using the `digital-signature-scheme` option at the [`edit security ike proposal` *`proposal-name`*] hierarchy level:

- `rsassa-pkcs1-v1-5` for RSA signature with the PKCS1 v1.5 scheme
- `rsassa-pss` for RSA signature with Probabilistic Signature Scheme (PSS) as the default signature scheme for RSA certificate

Do not configure `digital-signature-scheme` without setting `digital-signature` as the authentication method.

[See digital-signature-scheme, proposal (Security IKE), and show security ike security-associations.]

Support for trusted CA group for certificate validation in IKE with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Junos OS uses the trusted certification authority (CA) group in an IKE policy to validate peer certificates during IKE authentication. You can use the feature on a device that runs the iked process for the IPsec VPN service. Configure the trusted CA group to specify which CAs the device trusts for incoming peer certificate validation. Use `set security ike policy` *`policy-name`* `certificate trusted-ca trusted-ca-group` *`trusted-ca-group-name`* to reference the trusted CA group in an IKE policy for certificate validation.

[See certificate and IKE Policy with Trusted CA Group.]

What's Changed

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

Known Limitations

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

Open Issues

IN THIS SECTION

Learn about open issues in this release for cSRX.

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

Intrusion Detection and Prevention (IDP)

- For cSRX/ARM platform, the below three attacks need to be configured outside the attack groups for proper detection. SHELLCODE:X86:AVD-UTF8TLWR-STC SHELLCODE:PHP:BASE64-STC HTTP:NNMRPTCONFIG-EXE-RCE

  PR1853515

Resolved Issues

There are no resolved issues in this release for cSRX.

Junos OS Release Notes for EX Series

IN THIS SECTION

What's New

IN THIS SECTION

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

EVPN

Auto-detect UAP interfaces for Group-Based Policy (EX4000, EX4100, EX4400, EX4650, and QFX5120)—When configuring unified access policy, you no longer need to specify the UAP interface role. You can

configure the switch to auto-detect whether the interface is connected to another switch or to a Mist AP.

Dynamic activation of VLANs with 802.1X in an EVPN-VXLAN network (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, 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 support dynamic activation of VLANs with 802.1X in an EVPN-VXLAN network. With this feature, VLANs are initially created locally on the Junos device in a deactivated (inactive) state. When a user client successful authenticates with 802.1X, the authentication server assigns the VLAN to the local port. Junos then activates the VLAN and installs the host routes in hardware.

To enable dynamic activation of VLANs with 802.1X, use `set vlans` *`vlan-name`* `activate-on-local-ports` or `set routing-instances` *`routing-instance-name`* `vlans` *`vlan-name`* `activate-on-local-port`.

[See Dynamic Activation of VLANs with 802.1X.]

EVPN maintenance mode scheduling (EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, QFX5120-32C, QFX5120-48T, and QFX5120-48Y)—You can schedule the EVPN maintenance mode CLI to minimize traffic loss and streamline upgrades across `erb-leaf`, `core-distribution`, `collapsed-core`, and `ip-clos` deployments. Use `set protocols evpn maintenance-mode (collapsed-core | core-distribution | erb-leaf | ip-clos) action-type (force | validate-and-set) schedule-start-at` *`time`* `[schedule-end-at` *`time`*`]` to trigger timed designated forwarder (DF) preference changes, network isolation (NISO), and Type 5 route withdrawal. Enable the NTP-based DF with `set protocols evpn df-election-ntp`. The `force` option overrides precheck results. Validate with `request evpn validate-maintenance-mode-pre-checks` *`deployment-mode`* and monitor with `show evpn maintenance-mode-status`. For Data Center Interconnect (DCI) gateways, delay interconnect Ethernet segment identifier (I-ESI) recovery using `set protocols evpn interconnect-esi-bringup-delay` *`seconds`*.

[See EVPN Maintenance Mode for Multihomed Leaf Isolation.]

EZ-LAG simplified CLI settings for native VLAN ID and access mode interfaces, and trunk mode interface configuration change (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, 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)—You can use the following new statements to define the peer provider edge (PE) device interfaces to the servers in an easy EVPN LAG (EZ-LAG) simplified configuration:

- Specify a user-defined native VLAN ID for trunk mode interfaces using the `native-vlan-id` *`vlan-id`* statement at the `[edit services evpn server` *`server-name`*`]` hierarchy level. You must include the existing `vlan-id-list [`*`vlan-id-list ...`*`]` statement to generate a trunk mode interface configuration.

- Generate PE-to-server interface configurations in access mode instead of trunk mode using the `access-vlan-id` *`access-vlan-id`* statement instead of the `vlan-id-list` statement. You can set `access-vlan-id`at the `[edit services evpn server` *`server-name`*`]` or `[edit services evpn server` *`server-name`* `mac-vrf-instance` *`instance-id`*`]` hierarchy levels.

Also, the EZ-LAG commit script now generates trunk interface configurations with the `flexible-vlan-tagging` option instead of the `vlan-tagging` option.

[See https://www.juniper.net/documentation/us/en/software/junos/evpn/topics/topic-map/easy-evpn-lag-config-script#concept_example-native-vlan-access-mode.html, server (Easy EVPN LAG Configuration), and mac-vrf-instance (Easy EVPN LAG Configuration).]

Global AS override for auto-derived route targets in EVPN (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, 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)—EVPN now supports global autonomous system (AS) override for auto-derived route targets. You can configure independent AS control for import and export route target generation by using the **import-as** and **export-as** statements. Alternatively, you can configure unified AS control for both import and export route target generation by using the **as-num** statement. This feature provides additional control over the AS value used for EVPN auto-derived route target generation in multi-AS deployments.

[See Global AS Override for Auto-Derived Route Targets.]

High Availability (HA) and Resiliency

Standalone ISSU (EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, and EX4400-48T)—You can use standalone in-service software upgrade (ISSU) on a single Routing Engine switch to upgrade Junos OS with minimal data traffic loss. The upgrade maintains forwarding state for L2 or L3 forwarding, LACP, DHCP, 802.1X sessions, and firewall filters, ensuring service continuity during a brief control-plane interruption. The upgrade also preserves uninterrupted Power over Ethernet (PoE), ensuring service continuity for powered devices.

[See Understanding Standalone ISSU.]

Junos Telemetry

Support for genstate YANG data models (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204,

EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Junos telemetry supports the genstate YANG data model, a YANG-based model autogenerated from operational state data. Genstate models expose a subset of `show` command outputs through gNMI subscriptions. This feature allows a gNMI telemetry collector to subscribe directly to specific resource paths in the genstate model to retrieve the required state data. The genstate models mirror selected hierarchies and fields from Junos `show` commands, allowing you to query information such as interface attributes, protocol status, system metrics, and more.

For example, you can retrieve the `evpn-instance` information by configuring the gNMI collector to subscribe to the following resource path:

```
/genstate/evpn-instance-information/evpn-instance
```

In this path, `evpn-instance` is a field displayed in the output of the `show evpn instance` command.

Genstate resource paths are prefixed with `/genstate/`. The genstate YANG files are published on Github. See Genstate YANG Data Models to view the supported show commands and the root paths.

[See Model-Driven Telemetry, Data Models, and Explore Sensor Paths.]

Decoupled rendering for telemetry data streaming (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—The Junos telemetry infrastructure now separates data rendering from core application processing, enhancing system performance and scalability. Dedicated rendering applications handle telemetry data rendering, enabling originating applications to focus on primary functions. This decoupling of functions improves efficiency, particularly during high-volume telemetry data streaming. Users may notice changes in telemetry producer names, as rendering is performed by the corresponding render applications. No configuration changes are needed, and existing subscriptions remain unaffected. This enhancement aligns the telemetry framework with the updated architecture, optimizing

performance and scalability for telemetry data processing and streaming. See, Decoupled Rendering of Telemetry Data for a complete list of root sensor paths and respective new producer application names.

Interface statistics and error counter telemetry sensors for Mist integration (EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-24F, EX4100-H-24MP, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, and EX4400-48XP)—Junos Telemetry now supports sensors for error counters to meet Mist cloud integration requirements. This feature enables streaming of comprehensive error counter statistics, providing end-to-end visibility into interface performance and errors, and enabling detailed diagnostics in Mist-managed environments. Subscribe to the root resource path */state/interfaces/interface/counters/errors* to stream interface input and output error counters. View all supported sensors in the Junos YANG Data Model Explorer.

Junos XML API and Scripting

Event policy archive sites support specifying a routing instance (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When you configure an archive site for an event policy destination, you can specify the routing instance that the device uses to connect to the site. To configure the routing instance, include the `routing-instance` option at the `[edit event-options destinations` *`destination-name`* `archive-sites` *`url`*`]` hierarchy level.

[See Event Policy File Archiving and archive-sites (Event Policy).]

MAC Learning

Disable MAC table aging per VLAN (EX3400, EX3400-VC, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48P, and EX4400-48T)—On EX Series switches with support for the Enhanced Layer 2 Software (ELS) configuration style, you can disable MAC address aging per VLAN so that MAC entries in that VLAN do not age out. This configuration reduces flooding and traffic loss in VLANs with

stable endpoints and supplements the global MAC aging timer (default 300 seconds) or its adjustment. To disable MAC address aging, configure the `no-mac-aging` option at the VLAN level through the CLI.

[See switch-options (VLANs) and no-mac-aging.]

Support to learn the MAC-IP information of the host (EX3400)—You can use MAC-IP snooping for non-EVPN VLANs to learn the MAC-IP information of the host. To enable this feature, configure the `global-mac-ip-snooping` command in the `[edit protocols l2-learning]` hierarchy to enable this feature. The `show ethernet-switching mac-ip-table` command output displays the MAC-IP information. You can verify whether MAC-IP snooping is enabled for a VLAN by using the `show vlans extensive` command. You can disable the feature for a VLAN by using the new command `no-mac-ip-snooping` in the `[edit vlans <vlan-name> switch-options]` hierarchy.

[See global-mac-ip-snooping and no-mac-ip-snooping.]

Routing Protocols

Redistribution of IPv4 routes over IPv6 routes into BGP (EX4400-48MP)—Devices running Junos OS can forward IPv4 traffic over an IPv6-only network, which generally cannot forward IPv4 traffic. As described in RFC 5549, *Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop*, IPv4 traffic is tunneled from customer premises equipment (CPE) devices to IPv4-over-IPv6 gateways. These gateways are announced to CPE devices through anycast addresses. The gateway devices then create dynamic IPv4-over-IPv6 tunnels to remote CPE devices and advertise IPv4 aggregate routes to steer traffic. Route reflectors with programmable interfaces inject the tunnel information into the network. The route reflectors are connected through internal BGP (IBGP) to gateway routers, which advertise the IPv4 addresses of host routes with IPv6 addresses as the next hop. Currently, the dynamic IPv4-over-IPv6 tunnel feature does not support unified ISSU.

To configure a dynamic IPv4-over-IPv6 tunnel, include the `dynamic-tunnels` statement at the `[edit routing-options]` hierarchy level.

[See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.]

Software Installation and Upgrade

In-band ZTP in EVPN multihoming deployments (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, 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)—In-band zero-touch provisioning (ZTP) now supports EVPN multihoming deployments that use Layer 2 integrated routing and bridging (IRB) interfaces. In EVPN multihoming topologies, downstream devices can establish redundant connections to multiple upstream devices during Day 0 provisioning. Upstream devices participate in an EVPN-VXLAN fabric to provide Layer 3 onboarding connectivity through DHCP relay services.

[See In-band Zero Touch Provisioning.]

User Access and Authentication

Enhancement: TLS Support for TACACS+ (TACACS+/TCP Secure Transport) (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—This feature enhances the existing TACACS+ over TCP (TACACS+/TCP) implementation by adding support for Transport Layer Security (TLS), enabling secure TACACS+ communication. With this update, the device-side TACACS+ client is extended to establish TLS sessions with TACACS+ servers, ensuring encryption of authentication, authorization, and accounting (AAA) traffic. The enhancement includes TLS handshake handling, server certificate validation, and secure session establishment, while maintaining backward compatibility with non-TLS configurations. This improves overall security by protecting sensitive data from interception and aligns TACACS+ communication with modern security standards.

Global and Per-Interface Session Limit Control (AND EX4650, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, and EX4400-48XP)—This feature introduces configurable global and per-interface session limits for 802.1X to protect switch resources and ensure predictable scaling. Administrators can define these limits to control the number of concurrent authentication sessions and mitigate the impact of authentication surges. By default, each interface supports up to 100 sessions, with no global session limit enforced. When the configured session limits are exceeded, the switch rejects additional authentication attempts and generates corresponding syslog events, helping maintain system stability and providing visibility into limit violations.

IP Device Tracking for 802.1X in Non-DHCP Environments (EX3400) (EX3400 and EX3400-VC)—This feature introduces IP Device Tracking on EX3400 platforms to monitor and report host IP addresses on 802.1X-authenticated ports in environments where DHCP snooping is not available. It enables visibility and control for statically addressed clients by leveraging Layer 2 MAC-IP snooping events to dynamically learn MAC-IP bindings. The learned IP address is included as the Framed-IP-Address attribute in RADIUS interim accounting updates, allowing administrators to enforce IP-based policies on the RADIUS server. This enhancement improves client tracking and policy enforcement in non-DHCP deployments and supports secure, policy-driven network access.

IP Device Tracking for Connected Hosts in Non-DHCP Environments (AND EX4650-48Y-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, and EX4650)—This feature enables IP Device Tracking to monitor and report host IP addresses on 802.1X-authenticated ports in environments where DHCP snooping is not deployed. It provides enhanced visibility and control for statically addressed clients by allowing administrators to enforce IP-based policies via a RADIUS server. The 802.1X control plane dynamically learns MAC-IP bindings from Layer 2 MAC-IP snooping events and includes the learned IP address as the *Framed-IP-Address* attribute in interim accounting updates. This capability improves endpoint tracking and supports more granular policy enforcement in non-DHCP network environments.

Sequential DNS fallback for outbound SSH (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When a configured hostname resolves to multiple IP addresses via DNS, Junos OS automatically tries each address in sequence, regardless of IP family, until a connection succeeds or all addresses have been attempted. The `outbound-ssh` retry mechanism activates after all resolved addresses have been tried.

[See outbound-ssh.]

Support for OSPFv2 HMAC SHA-2 keychain authentication (EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, QFX5120-48Y, and QFX5120-48YM)—You can enable the OSPFv2 keychain module with HMAC-SHA2 authentication to authenticate packets reaching or originating from an OSPF interface. HMAC SHA2 algorithms include HMAC-SHA2-256, HMAC-SHA2-384, and HMAC-SHA2-512 as defined in *RFC 5709*. We also support these algorithms along with the HMAC-SHA2-224 algorithm. This feature ensures smooth transition from one key to another for OSPFv2 with enhanced security. We also support HMAC-SHA1 and HMAC-SHA2 authentication for virtual and sham links.

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

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

[See Understanding OSPFv2 Authentication.]

Additional Features

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

BGP Auto-discovery underlay in EVPN-VXLAN (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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, QFX5120-48Y, and QFX5120-48Y-VC)—We support BGP auto-discovery underlay in EVPN-VXLAN networks.

[See BGP Auto-Discovered Neighbors and peer-auto-discovery.]

L2 multicast enhancement to support Professional AudioVisual (ProAV) traffic (EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, and EX4000-48T)—We have enhanced L2 multicast to support ProAV traffic over IP networks with optimized bandwidth utilization.

[See Restrict Flooding of Unknown Multicast Traffic to the Multicast-router Interface in an L2 environment, flood-reports-to-inter-switch-interface (MLD Snooping), and immediate-leave-on-host-interface (MLD Snooping).]

Overlay and CE-IP ping and traceroute support with IPv6 underlay in EVPN-VXLAN (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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)—Overlay and CE-IP ping and traceroute support with IPv6 underlay in EVPN-VXLAN

[See ping overlay, ping, traceroute, and traceroute overlay.]

Support for hardware resource threshold monitoring for capacity planning (EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, QFX5120-48T-VC, and QFX5120-48Y-VC)—Hardware Resource Monitoring alerts you when a resource reaches an upper or lower threshold. This feature helps prevent traffic loss or device downtime through timely action.

[See Configure Hardware Resource Threshold Monitoring for Capacity Planning.]

Support for loopback firewall filter optimization and applying firewall filters to loopback address logical interfaces on a per VRF basis— (AND EX4400-48T, EX4100-24MP, EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, and EX4400-48P)—Support is added for the following CLI configuration commands:

- `set chassis loopback-firewall-optimization`
- `set chassis loopback-firewall-per-vrf`

[See loopback-firewall-optimization and loopback-firewall-per-vrf.]

Support for PKI (EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, EX4100- F-12T, EX4100-24MP, EX4100-24P, EX4100-24T, EX4100-48MP, EX4100-48P, EX4100-48T, EX4100-F-12P, EX4100-F-24P, EX4100-F-24T, EX4100-F-48P, EX4100-F-48T, EX4100-H-12MP, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, and EX4400-48XP)—We support PKI.

[See Public Key Infrastructure User Guide.]

Support for Quantum Buffer in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use

Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptographic agility.

[See The Quantum Buffer and moduli.]

Support for Shor-resistant and other default key exchange algorithms in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSH supports the Shor's algorithm-resistant hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm. And, by default, SSH supports the `dh-group16-sha512` and `dh-group18-sha512` Diffie-Hellman (DH) group key exchange algorithms.

[See key-exchange.]

Support for ZTP with IPv6 (AND EX9214, EX4000-12MP, EX4000-24MP, EX4000-48MP, EX4000-8P, 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-DC, EX4100-H-24F-DC, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, and EX9208)—Use a DHCPv6 client and zero-touch provisioning (ZTP) to provision a device.

[See Zero Touch Provisioning.]

Telemetry streaming for secure packet capture (EX4650, EX4650-48Y-VC, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)—This feature enables you to capture packets from a device and securely stream them to an external telemetry data collector (for example, a cloud-based collector) for monitoring and analysis. Network professionals use real-time packet capture data to troubleshoot complex issues, including network and performance degradation, and poor end-user experience.

For secure packet capture, the maximum packet size is 512 bytes, including both the packet header and payload. To enable secure packet capture, include the */junos/system/linecard/packet-capture* resource path in a Junos RPC call.

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

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

This configuration sends packet capture sensor data to the collector. Remove the `packet-capture` configuration after the required data is collected. After capture is complete, ingress packets that match the filter conditions are trapped to the CPU and forwarded to the collector over a secure channel using the JTI-defined key-value format over a gRPC connection.

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 configure multiple interfaces for egress packet capture. When enabled, host-bound egress packets are captured and sent to the collector. Remove the configuration when packet capture is no longer required.

[See Junos YANG Data Model Explorer and Supported gRPC and gNMI Sensors.]

What's Changed

IN THIS SECTION

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

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

EVPN

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit protocols evpn]` hierarchy level only on platforms that support a default switch instance. To enable these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the `[edit routing-instances name protocols evpn]` hierarchy level.

  [See Preference-Based DF Election.]

- **XML output tags changed for `show l2-learning evpn mcnh-info` command**—We've changed the XML output for this command. The XML `<l2ald-mcnh-entry-list-headers>` element level under the `<l2ald-mcnh-entry>` element has been removed. The child output tags that were under the `<l2ald-mcnh-entry-list-headers>` element are now direct child tags of the `<l2ald-mcnh-entry-list>` element.

General Routing

- **vxlan vni match condition enabled (EX4400-24MP, QFX5120-32C, QFX5120, QFX5120-32C)**—The following match conditions are now enabled on these platforms:

  - vxlan vni <vni-value>—Match the VNI.
  - vxlan vni-except <vni-value>—Do not match the VNI.

  [See Firewall Filter Match Conditions and Actions (PTX Series Routers).]

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the [`edit services port-extender satellite <name> device-model`] hierarchy, you can use the '?' to see a list of supported platforms. We have also added the `skip-lo0-config` statement under [`edit service port-extender jnu`] hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added `cp-mtu` under the [`edit service port-extender satellite <name>`] hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

  [See port-extender (Services).]

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for auto-configuration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address `ff02::1`. The RA source is the router's link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- Topic updated with Platform Specific Behavior section - https://uat-aem.juniper.net/documentation/test/us/en/software/junos/routing-policy/topics/concept/firewall-filter-flexible-match-conditions-overview.html

- **Enhanced loop detection for Layer 2 networks (EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)** —We've added support for monitoring all VLANs on a logical interface. You can set `vlan-id all` option under the [`edit vlans *vlan-name*`] hierarchy. With this change, the devices can detect network loops across multiple VLANs and interfaces, improving network stability and performance.

  [See Enable Lightweight PE-CE Loop Detect on a Logical Interface and loop-detect.]

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the `show route extensive expanded-nh` command.

  [See show route extensive <route> expanded-nh.]

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)               - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)        -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- **Manual reboot required for specific forwarding and chassis options in clustered deployments (EX Series, QFX Series)**—For a Virtual Chassis or Virtual Chassis Fabric (VCF), the Packet Forwarding Engine in member switches does not automatically restart upon configuring and committing. When you commit any of the following statements, the commit succeeds but the new profile does not take effect until you manually reboot the system: `set forwarding-options vrf-fallback set forwarding-options mpls-tunnel-extend set forwarding-options vxlan-disable-copy-tos-encap set forwarding-options vxlan-disable-copy-tos-decap set chassis loopback-firewall-optimization set chassis per-logical-interface-firewall set chassis loopback-firewall-per-vrf` After committing these configuration statements, you must schedule and perform a manual reboot for the changes to be applied to the forwarding plane on all members. Until you reboot, the system continues to use the existing forwarding profile.

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- Support added for new payload match conditions (EX4100 , EX4400 , 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 https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/firewall-filter-match-conditions-for-ipv4-traffic.html]

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **Per-fan entries for fan trays in SNMP MIB tables (PTX10008)**—You can monitor individual fans in a fan tray through distinct entries in SNMP MIB tables. This enhancement addresses the earlier limitation that SNMP MIB does not list individual fans.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.

  - A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Member Link Limit on Adaptive Aggregated Ethernet (PTX Series)**—A maximum of 64 member links are configurable on an aggregated Ethernet (AE) bundle with adaptive load balancing. Configuring more than 64 member links will result in a commit failure.

  [See Understanding Aggregated Ethernet Load Balancing.]

- Write cache disabled alarms no longer raised—The "Disk 2 Write Cache disabled" alarms were reported in the chassis alarms due to a missing SSD model in Junos OS supported device table. The support is now added, and alarms are no longer reported.

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading,

copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

[See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:
  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.
  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.
  - The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

  These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

  [See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Network Management and Monitoring

- **Improved ZTP SNMP configuration handling (Junos)**—We have reordered the SNMP configuration commands to optimize the Zero Touch Provisioning (ZTP) process. The `set snmp engine-id local` statement now processes before `set snmp v3 usm local-engine user`, ensuring SNMPv3 USM authentication and privacy keys generation using the configured local engine ID instead of an auto-derived default. This prevents SNMPv3 query authentication failures from engine-ID resets.

  [See Configure Engine ID on SNMPv3 and Create SNMPv3 Users.]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for

each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

[See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **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.]

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **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.

User Interface and Configuration

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:

- Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`
- Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`
- Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

[See show system data-models.]

- **Introduction of remove operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents <userinput>statement not found</userinput> warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

- **CLI options and RPC formats to retrieve `genstate` information and XPath expressions for operational commands**—A gNMI telemetry collector can subscribe to `genstate` resource paths to query for the corresponding state data on a device. We've added CLI pipe filter functions and RPC formats that enable you to discover available `genstate` resource paths and the associated `genstate` YANG module for a given show command. You can also retrieve the XPath expression for any node in command output. The options include:

  - `| display genstate subscription-paths`—Display all resource paths available for subscription in the `genstate` YANG module associated with the given command.
  - `| display genstate yang-module-name`—Display the name of the `genstate` YANG module that includes the data model and resource paths for the given command.
  - `| display xml genstate-path`—Display the command output in XML and include the `junos:genstate-path` attribute for every emitted tag that is available for subscription. The attribute value is the `genstate` subscription path. For RPCs, use the `format='xml-genstate-path'` attribute in the request tag.
  - `| display xml rpc response-paths`—List all possible response paths under the equivalent RPC's top-level response tag.
  - `| display xml xpath`—Display the command output in XML and include the `junos:data-xpath` attribute for each tag in the output. The attribute value is the XPath expression that selects that specific node. For RPCs, use the `format='xml-xpath'` attribute in the request tag.

[See | (pipe) and Junos Genstate YANG Data Models.]

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

Known Limitations

IN THIS SECTION

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

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

General Routing

- With SFP-T optics and DUT configured speed=100M with Auto-Neg enabled and peer configured speed=10M with Auto-Neg enabled. Link will not come up PR1905956

Routing Protocols

- larger primes of size greater than 8192 are generated when designer type moduli is confiugured with count 3 and 4

  PR1942246

Open Issues

IN THIS SECTION

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

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

General Routing

- 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 may 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
- When a interface supports multiple lanes, the SNMP OID jnxDomCurrentLaneWarnings is incorrectly handled as a single lane, resulting in the value from lane 0 being replicated across all other lanes. PR1906927
- request system firmware upgrade jfirmware allfw only upgrading only one of the components even when multiple components are available for upgrade PR1942521
- EX: OIR of PSU within short intervals can affect APs not to come up in N+1 PSU redundancy mode. Work-Around: restart chassisd daemon helps to recover from the issue. Recommendation: Please wait enough (~30s) between OIRs of PSUs PR1939315
- On EX4400 platforms, the chassis daemon (chassisd) crashes when the configuration statement `set chassis alarm fru-absence ignore fpc <slot> pem <slot>` is committed.

  PR1956927

Platform and Infrastructure

- 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

Resolved Issues

IN THIS SECTION

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

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

EVPN

- When Junos and Junos Evolved platform sends two proxied Neighbor Advertisement for default gateway, second one without Router flag and target Link-Layer address. The second Neighbor Advertisement can disrupt connectivity for end hosts. A proxied Neighbor Advertisement (NA) is an IPv6 mechanism where a network device responds to Neighbor Discovery requests on behalf of another device, instead of the actual endpoint replying itself. PR1926468

General Routing

- On Junos EX4100 and EX4400 platforms, inline-monitoring for layer 2 and layer 3 will face issue due to hardware limitation. There needs to be separate FTFP groups in hardware Layer 2 and Layer 3 inline monitoring. It will not work simultaneously for L2 and L3 inline-monitoring. Only one group will remain active and monitoring will work for the first group. PR1926180
- On all EX platforms during the shutdown regular sequence a crash could occur in rare scenarios generating a core log file. PR1932731

- On EX4000 and EX4100 platforms, during zeroization, factory installed identity certificate (iDevID) is deleted, which impacts applications that depend on the iDevID (e.g. sZTP)

  PR1895474

- On all Junos OS platforms and Junos OS Evolved platforms which supports storm control, when a different storm-control profile is applied for interface where these profile have same profile index allocated then the storm control profile configuration and system state will not be in sync and a different profile will be applied for interface binding due to profile index collision which results into l2ald process crash. PR1930380

- On all Junos EX Series platforms, following a Virtual Chassis (VC) reboot, random interface links GE(Gigabit) and XE(10 Gigabit) intermittently appear in a down state at the Interface Descriptor (IFD) level, even though the physical link remains operational.

  PR1915222

- In an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) deployment, a remote MAC move event can cause the remote MAC IP entry to be removed after it is initially installed, resulting in missing Address Resolution Protocol (ARP) state and incomplete hardware programming on the remote Virtual Tunnel Endpoint (VTEP), which can lead to traffic loss for the affected endpoint. PR1936648

- On EX4400 platforms with 4x25G or 1x100G ULM (Universal Link Module), the CPU hogs by CMQFX thread for as much as 3.7 secs when Firmware download occurs during PIC offline/online and PFE restart time / Device reboot. To speed up the firmware download operation, the MDIO clock (MDC) frequency is increased from 2.6 MHz to 9 MHz which reduced firmware download time from ~3.8 s to ~2.2 s. This is an enhancement PR. PR1870962

- On EX4400 platforms, due to an internal communication issue within the CPU controller, the system triggers repeated fxpc panics which in turn lead to the dcpfe process (Dataplane Packet Forwarding Engine) to crash and restart continuously, resulting in service disruption.

  PR1919727

- On all Junos devices with dual RE configured with dot1x, if the device is upgraded from release 23.4R2-S4 or there is a RE switchover and persistent-cache feature enabled it will cause dot1xd to crash. It is due to presence of some stale files on backup node. PR1923266

- On Junos EX2300 and EX3400 switches, after the device or the Packet Forwarding Engine (PFE) starts, some hardware telemetry or statistics related to Network Processing Unit can appear incorrect or fail to report as expected; network traffic is not affected.

  PR1922040

- On EX4100 platform VCP link flap due to SYSPLD read failures and mark SFP as unplugged PR1904884

- On Junos EX and QFX platforms, with Virtual Extensible Local Area Network (VXLAN) configured, during reboot, power cycle or upgrade, upon completion of the VXLAN initialization process, when the VXLAN BUM flag is disabled by config "set forwarding-options vxlan-bum-bit-disable, the parameter returns to its default value and hence the vxlan bum bit cli configuration doesn't work as expected. PR1925572

- On Junos-based EX and QFX platforms, modifying any term of a firewall filter applied on the lo0 interface triggers repeated Dynamic Firewall Daemon (dfwd) restarts and causes the Routing Engine (RE) CPU to reach 100 percent. The condition persists until the filter is removed or the previous configuration is restored. PR1925850

- IPv6 neighbor is not switching to the correct Redundant Trunk Group (RTG) link when we bring back up the primary link after a failover scenerio. Switching from the primary link to the standby works ok, but when switching back, the ipv6 neighbor still points to the standby link causing traffic to blackhole to that destination. Ipv4 arp works ok. PR1939656

- On all Junos EX and QFX platforms operating in an Ethernet Virtual Private Network (EVPN) environment, this issue occurs when the CLI command `licensing hourly update` is executed on a device that does not support the licensing feature.

  PR1930036

- On Junos OS EX series Virtual Chassis platforms, the MACsec interface fails to re-establish sessions after multiple reboots of a particular Virtual Chassis (VC) member or when a VC link flaps. This issue is observed between 4/5 reboots of the VC member. In these case, the MACsec daemon is not notified that the license is installed on the member, resulting in MACsec encryption being unavailable on that interface.

  PR1914499

- On EX4400 platforms, after PIC online/Soft OIR of optic/Optic removal and insertion/Reboot, 'sfp_is_non_ieee_jnpr_sfpp Unsupported optic, not a 25G Bidi optic Custom Opt 0 0x6 custom opt' message is seen.

  PR1940241

- On Junos platforms having QFX and EX, VLAN(Virtual Local Area Network) traffic loss occur on interfaces configured with flexible-vlan-tagging when an L3 (SP style) sub-interface is deactivated on a port that also carries L2 (enterprise style) VLAN units. PR1938291

- Intermittent uplink ports flaps events observed and log messages show continuous syspld read failures during the issue. After a power cycle, the devices recovered. PR1925996

- On Junos QFX5110, QFX5120, EX4650, EX4400, EX4100 platforms, when operating with EVPN-VxLAN type5 overlay ECMP at high tunnel next hop scale, the Packet Forwarding Engine (PFE) crash when a memory corruption issue is observed due to a buffer overflow that leads to corruption of

heap management structures. This corrupted metadata is detected by the memory manager when later heap allocation or free operations are invoked, resulting in a reported heap corruption condition.

PR1933698

- jnxFruName and jnxFruState are invalid values when snmp mib gets executed. > show chassis environment Class Item Status Measurement Power FPC 0 Power Supply 0 OK 34 degrees C / 93 degrees F FPC 0 Power Supply 1 Absent <<<<<<<<<<<<<<<<< > show snmp mib walk jnxFruName jnxFruName.2.1.1.0 = Power Supply 0 @ 0/0/* <<<<<<<< only PS0 jnxFruName.4.1.1.0 = Fan Tray 0 @ 0/0/* jnxFruName.4.1.2.0 = Fan Tray 1 @ 0/1/* --snip-- > show snmp mib walk jnxFruState jnxFruState.2.1.1.0 = 6 <<<<<<<< only PS0 jnxFruState.4.1.1.0 = 6 jnxFruState.4.1.2.0 = 6 --snip-- PR1934571
- Intermittent performance issues observed across sites due to anomalous DHCP behavior, where DHCP server packets are incorrectly received on downstream switch interfaces, triggering snooping drops. PR1908971
- On EX4400 platforms, loss of PoE (Power over Ethernet) power is seen on ports after port bounce or even during normal operation without specific external trigger . This causes outage on all PoE ports and devices connected to PoE port will not receive power causing service impact. PR1930080
- On platforms supporting dot1x authentication, when a Radius accounting Interim-Update is manually triggered from the switch (using a COA), the resulting accounting packet does not include the Framed-IP-Address attribute. However, when the same session generates an interim update via the periodic timer (10-minute interval), the Framed-IP-Address is included as expected. PR1934680
- On EX4400 product family platforms, device experienced an unexpected reboot and recovers automatically. The system logs indicate a kernel panic, and a VMcore file is generated. Traffic impact will be seen until the recovery. Bios version CDEN_P_EX1_00.22.01.00 fixes this issue. PR1942085
- On QFX5120, QFX5200, EX4400 platforms running Junos, when a firewall filter term includes the then sampling action and the user configures the sFlow sample rate only at the interface level using set protocols sflow interfaces <interface> sample-rate ingress <value>, the commit operation fails with the error: "sampling action required sflow sample-rate ingress and the sflow interface configured under protocols" PR1942162
- On Junos OS QFX5k and EX4k platforms, executing the VTY command show firewall filter hw filter-name <non-existent-name> when a firewall filter is configured triggers a crash of the dcpfe process. This occurs when the command references a hardware filter name that does not exist. This results in a PFE (Packet Forwarding Engine) restart, leading to temporary forwarding plane traffic disruption while the process recovers.

  PR1944897

J-Web

- When a MAC-RADIUS authentication succeeds and the RADIUS server returns CWA redirect attributes (URL-Redirect), the EX switch should intercept client HTTP traffic and return an HTTP 302 redirect to the Captive Portal. The switch instead responded with HTTP 405 (Method Not Allowed), causing the client to bypass the CWA redirect workflow. PR1931780

Layer 2 Ethernet Services

- On all Junos and Junos Evolved platforms, when a loop occurs in the transmission switch, the device starts receiving looped LACP (Link Aggregation Control Protocol) PDU's from itself, instead of messages from the actual peer device. This causes the system to mistakenly believe that a valid LACP connection exists, even though the peer device is not actually connected.As a result, it continues to forward traffic as if the peer were active. Since no valid peer connection is present, this can lead to traffic blackholing . PR1874126

Routing Protocols

- On all Junos Platforms, a rare memory corruption condition may occur in the Packet Forwarding Engine (PFE) when Spanning Tree Protocol (STP) is enabled and operating in default (distributed) mode. When the issue is triggered, the PFE crashes and a dc-pfe core file is generated. The exact trigger for the memory corruption is currently unknown. PR1931633

Junos OS Release Notes for JRR Series

IN THIS SECTION

What's New

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

What's Changed

IN THIS SECTION

Learn about what's changed in this release for JRR Series.

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

EVPN

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit`

protocols evpn] hierarchy level only on platforms that support a default switch instance. To enable these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the [edit routing-instances name protocols evpn] hierarchy level.

[See Preference-Based DF Election.]

General Routing

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for auto-configuration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address ff02::1. The RA source is the router's link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

Network Management and Monitoring

- **Improved Syslog User Identification (SRX4300)**—Syslogs generated from shell commands now logs the remote username instead of the local template username when you use the RADIUS or TACACS + authentication.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The show bgp neighborcommand output has been enhanced to display whether add-path send and receive capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and

additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **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.]

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

Known Limitations

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

Open Issues

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

Resolved Issues

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

Junos OS Release Notes for MX Series

IN THIS SECTION

What's New

IN THIS SECTION

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

Hardware

New MX301 Router Model (MX301-M)—A new model (MX301-M) of the MX301 Universal Routing Platform is now available with 64-GB of system memory. The MX301-M is designed for deployments with more moderate routing and service scale requirements, providing customers with an additional platform option. The MX301-M maintains the same hardware design, interfaces, forwarding capabilities, and software features as the MX301, while providing greater deployment flexibility across a broader range of network scale requirements.

To install MX301-M and perform initial configuration, routine maintenance, and troubleshooting, see the MX301 and MX301-M Universal Routing Platform Hardware Guide.

For a detailed list of available features, see Feature Explorer.

NOTE:
The MX301-M supports the same set of software features as the MX301. As such, they are listed together in Feature Explorer. Please note, however, that the MX301-M was first released in 26.2R1 and MX301 was first released in 25.4R1. Any 25.4R1 entries in Feature Explorer only apply to the MX301.

Dynamic Host Configuration Protocol

DHCP fingerprint forwarding support (DHCP local server and relay) (MX204, MX240, MX301, MX304, MX480, and MX960)—Enable DHCP fingerprinting forwarding to send validated client DHCP messages to fingerprint servers. This action preserves normal relay or local server processing. Fingerprinting servers analyze DHCP option sets and ordering to derive the device fingerprints. Use this capability to detect and block unauthorized devices, classify endpoints, assign VLANs or subnets by type, and apply device-specific policies. Configure `fingerprint-server-group` and `active-fingerprint-server-group` under `dhcp-local-server` or `forwarding-options dhcp-relay`. Optionally, set `logical-system`, `routing-instance`, and per-group overrides.

[See DHCP Fingerprint Forwarding Support (DHCP Local Server and DHCP Relay).]

EVPN

Headend termination of EVPN VPWS to EVPN E-LAN pseudowires with multihoming and flexible cross-connect interworking (MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10004, and MX10008)—Use pseudowire (ps) interfaces to terminate EVPN-VPWS pseudowires into EVPN-ELAN, streamline service stitching, and increase resiliency. You configure the ps transport interface for VPWS access and the ps service interfaces for Ethernet LAN (E-LAN). Apply ESI-based DF election in active/active or active/standby mode. Optionally use VLAN-unaware flexible cross-connect (FXC) to conserve labels. Transports include LDP, RSVP-TE, and SR-MPLS, with egress link protection, control word, and flow label options. Use existing operational commands for configuration and verification. EVPN E-LAN does not support VLAN-aware FXC termination.

[See Overview of EVPN VPWS into EVPN E-LAN via PS interfaces and PWHT.]

Support for using a non-default routing table for EVPN-VXLAN next-hop resolution (MX204, MX301, MX304, MX480, MX960, MX10003, MX10004, and MX10008)—We support the use of a non-default routing table to resolve next-hop IP addresses in an EVPN-VXLAN network, particularly when the network deployment spans multiple regions in a public cloud. In these scenarios, the next-hop IP addresses for the EVPN-VXLAN overlay network are not available in the default routing table (inet.0 or inet6.0) so Junos cannot resolve the next hop.

With this release, Junos now supports a forwarding context that maps BGP neighbors to `vrf` routing instances. By defining a resolution scheme and applying color-based policies, you can route EVPN-VXLAN traffic to use specific VRF routing table instances. Junos will then use the VRF routing table to resolve the next hop. This feature allows you to have more granular control when mapping VLAN traffic over traffic flows.

[See Non-Default Routing Table for EVPN Next-Hop Resolution.]

Forwarding and Sampling

Configuring Maximum ECMP Beyond the Default 128 Next Hops (MX301, MX304, MX480, MX960, MX2010, and MX2020)—You can use the `advanced-forwarding` mode to configure maximum ECMP beyond the default 128 next hops to up to 512 next hops on AFT-based line cards. This is useful for example when you have enabled symmetric consistent hash towards remote dynamic tunnel destinations (MPLS over UDP).

[See advanced-forwarding.]

Junos Telemetry

Support for genstate YANG data models (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T,

EX4000-8P, 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-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Junos telemetry supports the genstate YANG data model, a YANG-based model autogenerated from operational state data. Genstate models expose a subset of `show` command outputs through gNMI subscriptions. This feature allows a gNMI telemetry collector to subscribe directly to specific resource paths in the genstate model to retrieve the required state data. The genstate models mirror selected hierarchies and fields from Junos `show` commands, allowing you to query information such as interface attributes, protocol status, system metrics, and more.

For example, you can retrieve the `evpn-instance` information by configuring the gNMI collector to subscribe to the following resource path:

```
/genstate/evpn-instance-information/evpn-instance
```

In this path, `evpn-instance` is a field displayed in the output of the `show evpn instance` command.

Genstate resource paths are prefixed with `/genstate/`. The genstate YANG files are published on Github. See Genstate YANG Data Models to view the supported show commands and the root paths.

[See Model-Driven Telemetry, Data Models, and Explore Sensor Paths.]

Telemetry for BGP diagnostics via YANG model and sensors (MX2010)—You can stream BGP diagnostics telemetry using a dedicated YANG module to accelerate detection and remediation with less manual effort. Subscribe to sensors for global, group, neighbor, NLRI, and local RIB data through management interfaces such as NETCONF. Examples include session state, flap reasons, import and export evaluation triggers, and label and threshold metrics. You do not need to change the CLI. Certain recommendation and route fields emit static values.

[See Junos YANG Data Model Explorer.]

Addition of send-community-type and telemetry support (MX960 and MX10016)—Extend core functionality to add `send-community-type` and telemetry support.

The following paths are supported:

```
/network-instances/network-instance/protocols/protocol/bgp/global/afi-safis/afi-safi/config/send-community-type
```

```
/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/config/send-community-type
```

/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/config/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/config/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/afi-safis/afi-safi/config/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/global/afi-safis/afi-safi/state/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/state/send-community-type

/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/afi-safis/afi-safi/state/send-community-type

The OpenConfig community-type doesn't define IPv6-ext-community types, so they are not controlled through the OC configuration and they are always advertised.

[See Junos YANG Data Model Explorer.]

Configuration of administrative distance for BGP (MX10008 and MX10016)—You can configure administrative distance (preferences) separately (internal or external) for BGP learned routes and corresponding paths.

[See Junos YANG Data Model Explorer.]

Graceful restart for IS-IS and BGP via OpenConfig without native CLI dependency (MX204, MX240, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—Use OpenConfig to manage IS-IS and BGP graceful restart, reducing reliance on a native global option and simplifying vendor-neutral automation.

To configure, enable graceful restart and set a 30–300 second T3 timer with */network-instances/.../isis/global/graceful-restart/config/enabled* and *.../config/restart-time.*

Note: You do not need additional CLI commands. IS-IS graceful restart operates only in the default instance**.** Avoid partial protocol GR enabling. Native configurations take precedence in mixed mode.

[See Junos YANG Data Model Explorer.]

Decoupled rendering for telemetry data streaming (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—The Junos telemetry infrastructure now separates data rendering from core application processing, enhancing system performance and scalability. Dedicated rendering applications handle telemetry data rendering, enabling originating applications to focus on primary functions. This decoupling of functions improves efficiency, particularly during high-volume telemetry data streaming. Users may notice changes in telemetry producer names, as rendering is performed by the corresponding render applications. No configuration changes are needed, and existing subscriptions remain unaffected. This enhancement aligns the telemetry framework with the updated architecture, optimizing performance and scalability for telemetry data processing and streaming. See, Decoupled Rendering of Telemetry Data for a complete list of root sensor paths and respective new producer application names.

OpenConfig YANG model support for BGP (ACX5448, MX204, MX240, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, and vRR)—Junos telemetry now supports the OpenConfig YANG model for BGP, providing a vendor-neutral framework for configuration and monitoring BGP. Retrieve BGP-specific operational and state data by subscribing to sensors under resource path hierarchies such as `inline communities`, `community-set`, `match-community-set`, `match-ext-community-set`, `peer-group`, `neighbor`, `policy-definition`, `administrative distance`, and `as-path-set`. OpenConfig-defined default behavior for BGP policies is implemented, with `community-level` policy options supporting `any-match` and `invert-match` configurations.

Use the Junos YANG Data Model Explorer to view supported sensor paths and platforms. The supported OpenConfig version for the openconfig-network-instance model is 4.4.1.

The following table lists the supported YANG models and their versions:

BGP YANG Model	Version
openconfig-routing-policy.yang	3.4.2
openconfig-bgp-common-multiprotocol.yang	9.9.1
openconfig-bgp-common-structure.yang	9.9.1

(Continued)

BGP YANG Model	Version
openconfig-bgp-common.yang	9.9.1
openconfig-bgp-errors.yang	6.1.0
openconfig-bgp-global.yang	9.9.1
openconfig-bgp-neighbor.yang	9.9.1
openconfig-bgp-peer-group.yang	9.9.1
openconfig-bgp-policy.yang	8.1.0
openconfig-bgp-types.yang	6.1.0
openconfig-bgp.yang	9.9.1
openconfig-types.yang	1.0.0

[See Junos YANG Data Model Explorer.]

Update to BGP link bandwidth community value (MX10016)—Support BGP to allow regular expressions 0 and 2^32 in the link bandwidth community value. The OpenConfig path is

`/routing-policy/defined-sets/bgp-defined-sets/ext-community-sets/ext-community-set/ .`

[See Junos YANG Data Model Explorer.]

Enhancement of BGP XPath policy (ACX5448, MX480, MX960, MX2020, MX10004, and vRR)—Add OpenConfig BGP XPath enhancements for default policy and policy definitions.

[See Junos YANG Data Model Explorer.]

Update, extend, and provide options for OpenConfig telemetry (MX480 and MX960)—Support has been added for OpenConfig BGP telemetry in these three ways:

- BGP telemetry has been updated to the latest OpenConfig models.
- Support has been added for the following `set-med` paths :

`/routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/config/set-med-action`

`/routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/state/set-med-action`

`/routing-policy/policy-definitions/policy-definition/statements/statement/actions/bgp-actions/state/set-med`

- Support for OpenConfig `match-set` options in the routing policy (the `any-match` option is added to the Junos CLI). The following paths are implemented:

`/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-community-set/config/match-set-options`

`/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-ext-community-set/state/match-set-options`

`/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-ext-community-set/config/match-set-options`

`/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-community-set/state/community-set`

`/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-community-set/config/community-set`

`/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-ext-community-set/state/ext-community-set`

`/routing-policy/policy-definitions/policy-definition/statements/statement/conditions/bgp-conditions/match-ext-community-set/config/ext-community-set`

[See Junos YANG Data Model Explorer.]

Junos XML API and Scripting

Event policy archive sites support specifying a routing instance (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, SRX300, SRX320,

SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When you configure an archive site for an event policy destination, you can specify the routing instance that the device uses to connect to the site. To configure the routing instance, include the `routing-instance` option at the `[edit event-options destinations` *`destination-name`* `archive-sites` *`url`*`]` hierarchy level.

[See Event Policy File Archiving and archive-sites (Event Policy).]

MACsec

Support for a custom EAPoL EtherType to improve network tunneling of MACsec packets (MX240, MX301, MX304, MX480, MX960, MX10004, and MX10008)—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 the EAPoL EtherType, use the `ether-type` *`ether-type-value`* statement at the `[edit forwarding-options custom-eapol-ether-type-profiles` *`eapol-profile-name`*`]` hierarchy level. You must use an EtherType that isn't already reserved for another use. 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. To view the new EtherType profile, use the `show security mka sessions detail` command.

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

MPLS

RSVP-TE SRLG Hold-Down for CSPF Path Selection (MX480, MX960, MX10008, and MX10016)—Shared Risk Link Group (SRLG) hold-down mechanism allows Constrained Shortest Path First (CSPF) to avoid recently unstable network resources when computing RSVP-TE label-switched paths (LSPs). When a link associated with an SRLG experiences a link-down event, the SRLG enters a hold-down state for a configurable period. During the hold-down time, CSPF treats links belonging to that SRLG as undesirable by either penalizing links by adding a configurable SRLG cost to the traffic engineering (TE) metric or excluding those links from CSPF path computation. You can configure SRLG hold-down policy at the global level under the `[edit protocols mpls srlg-hold-down]` hierarchy level or for individual SRLG under the `[edit routing-options srlg` *`srlg-name`*`]` hierarchy level.

[See Shared Risk Link Groups for MPLS.]

Multicast

MVPN HRS enhancements (MX204, MX240, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—You can use the following configurations to enhance hot-root standby (HRS):

- To refine upstream multicast hop (UMH) selection by considering multicast LDP tunnel health, enable the `tunnel-status` option at the `[edit protocols mvpn umh-selection-additional-input]` hierarchy level. The device excludes ingress provider edge (PE) routers without valid tunnels and retriggers selection based on the egress label status.

- To revert MVPN flows to the stable primary path on the egress PE router, issue the `request mvpn hrs-force-revert` command. You can filter the reversion by instance, group, and source.

If a PE router is configured as a `receiver-only` site, the route-target import community is removed from advertised unicast routes to prevent UMH selection.

[See tunnel-status, request-mvpn-hrs-force-revert, and receiver-site.]

Network Management and Monitoring

Enhanced gr- interface count and avoidance of loopback stream in port mirroring (MX304, MX480, MX960, MX2010, and MX2020)—Support for port mirroring to sample the ingress or egress traffic of a `gr-` interface and avoid the loopback stream of the `gr-` interface. Port mirroring---using family `inet`, `inet6`, or `bridge`---samples packets ingressing or egressing a `gr-` interface and sends those samples to the mirror destination, where the packet analyzer is attached.

Port mirroring that belongs to the `bridge` family and that is applied in the output direction for a `gr-` interface might not work.

Caveats of local and remote port mirroring apply.

[See Configuring Port Mirroring.]

Juniper Routing Director support (MX301)—You can use Juniper Routing Director (formerly Juniper Paragon Automation) to onboard, provision, and manage the MX301 router. Juniper Routing Director provides end-to-end WAN automation, simplifying intent-based networking across the entire device life cycle from initial deployment and service provisioning to AI-driven optimization and active assurance.

[See Onboard, Configure, and Monitor MX301 and MX301-M.]

NextGen Port Extender (NGPE)

Centralized syslog management with structured, priority-based logging, UDP forwarding, and IPv6 support (MX304, MX10004, and MX10008)—Use centralized syslog configuration on an aggregation device to enforce uniform logging across satellite devices. Satellites forward logs over UDP to the aggregator, which collects and relays them to external syslog servers for comprehensive diagnostics. Configure structured and priority-based logging on satellites to enhance monitoring and analysis while maintaining security and compatibility, with IPv6 support.

Precision Time Protocol (PTP)

Precision Time Protocol with MACsec encryption (MX301)—On 1GbE, 10GbE, and 100GbE interfaces, you can use Precision Time Protocol (PTP) with Media Access Control Security (MACsec) encryption enabled on the same port. This configuration ensures secure and precise time synchronization, enhancing network performance and security.

[See G.8275.1 Telecom Profile, PTP with Media Access Control Security (MACsec) Encryption, and Guidelines to Configure PTP over Ethernet.]

Synchronous Ethernet support (MX301)—We support the following Synchronous Ethernet timing features on 1GbE, 10GbE, and 100GbE interfaces:

- Synchronous Ethernet over link aggregation group (LAG)
- Synchronous Ethernet over LAG with Ethernet Synchronization Message Channel (ESMC)
- Synchronous Ethernet over a MACsec-enabled port
- Synchronous Ethernet over LAG with enhanced ESMC. The enhanced Ethernet equipment clock (eEEC) supports new clocks with different quality levels in the Synchronous Ethernet chain and transmits ESMC with extended quality level TLV.

[See Synchronous Ethernet and enable-extended-ql-tlv.]

Routing Options

Route re-resolution optimization for L3VPN and recursive multipath with new diagnostic commands (MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—You can now use optimized route re-resolution in the resolver module to reduce CPU utilization during route churn. This optimization extends re-resolution beyond inet.0 and inet6.0 to `<instance>.inet.0`, `<instance>.inet6.0`, and L3VPN tables (bgp.l3vpn.0, bgp.l3vpn-inet6.0). The resolver also supports optimized re-resolution for recursive resolution cases and for resolution over multipath prefixes.

To monitor and diagnose resolution, you can use the following operational commands:

1. `show route resolution active-helpers` - Displays filtered routes used for the resolution, before the active routes were displayed.
2. `show route resolution active-helpers statistics` - Helps identify the reason for which routes are re-resolved, which is helpful in debugging.
3. `show route resolution re-resolution-history` - Displays the routes that are re-resolved.
4. `show route resolution loop-history` - Helps identify if any resolution loop occurred.

[See show route resolution.]

Routing Protocols

Support for handling unknown optional transitive BGP attributes (MX960 AND MX10008)—You can drop unknown optional transitive BGP attributes or withdraw routes carrying these attributes to prevent the propagation of malformed or harmful attributes across BGP peers. Apply the configurations at the BGP global, group, or neighbor level for comprehensive control. You can specify exceptions for attributes that should not be dropped or cause route withdrawal.

`Note:` If you configure multiple actions for the same attribute, then you can perform only the withdrawal action.

To display the unknown optional transitive BGP attributes, use the enhanced `show BGP neighbor` and `show log messages` commands.

[See show bgp neighbor.]

Services Applications

GRE Loopback-Avoidance Mode and GRE-Only Tunnel Services Scaling (MX301, MX304, MX480, MX960, MX2010, MX2020, MX10004, and MX10008)—By default, when a packet is encapsulated on a generic routing encapsulation (GRE) interface, it is sent through the loopback, where a second lookup is performed on the outer IP-GRE header to determine its forwarding path. With loopback avoidance, the packet bypasses this loopback stage and is forwarded immediately after GRE encapsulation in a single pass, preserving forwarding bandwidth and improving efficiency. Loopback avoidance on GRE logical interfaces is supported for the inet, inet6, and bridge families.

GRE logical interface (IFL) scale was limited to 4K per Physical Interface Card (PIC). Newer support increases this scale by dedicating tunnel resources on a PIC exclusively to GRE IFLs, allowing up to 32K GRE IFLs per chassis. Within that overall limit, a maximum of 16K IFLs is supported per interface device (IFD), whether it is a GRE IFD or a physical IFD.

Inline monitoring services to export sampled packets in sFlow version 5 format (MX304)—Inline monitoring services can now export `family inet` and `family inet6` sampled packets in sFlow version 5 format instead of the current IPFIX format. IPFIX format is the default. Configure the `export-format sflow-v5` statement at the `[edit inline-monitoring instance` *`instance-name`*`]` hierarchy level to change the export format. You do not need to configure sFlow (using the `set protocols sflow` command) to activate this feature.

[See instance (Inline Monitoring) and Inline Monitoring Services Configuration.]

Stateless MAP-T with VRF separation, DHCPv6 provisioning, inline IP reassembly, and per-flow ECMP (MX301, MX304, MX480, MX960, MX10004, and MX10008)—You can deliver stateless IPv4-IPv6 translation across an IPv6-only core using MAP-T while isolating subscriber and Internet traffic in VRF instances via next-hop inline services. Provision CPE MAP-T parameters dynamically via DHCPv6 option

95, enable inline IP reassembly for IPv4/IPv6 fragments, and apply per-flow ECMP across multiple service paths. This approach enhances scalability, resiliency, and traffic separation without relying on stateful translators.

Source Packet Routing in Networking (SPRING) or Segment Routing

BGP RIB sharding for SRv6 routes (MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—We've extended support for Segment Routing over IPv6 (SRv6) with BGP routing information base (RIB) sharding to enhance route processing and utilize SRv6 next hops. RIB sharding splits BGP service resolution over multiple shards to improve BGP convergence and performance. Configure SRv6 under IS-IS or as static SRv6-Traffic Engineering (SRv6-TE) with BGP RIB sharding. We do not support BGP route resolution when RIB sharding is configured.

[See rib-sharding.]

Computed SR-TE policy for IPv4 and IPv6 anycast and non-router ID endpoints in SR-MPLS (MX960, MX2010, and MX2020)—You can use IGP-learned non-router-ID destinations such as IPv4 and IPv6 anycast prefixes, secondary loopbacks, or interface prefixes as endpoints in SR-MPLS Segment Routing-Traffic Engineering (SR-TE) policies. Apply locally computed distributed Constrained Shortest Path First policies with optional compute profiles and segment identifier (SID) stack compression. The ingress computes candidate paths using the traffic engineering database (TED), metrics, and constraints, steering traffic to the nearest eligible anycast member. This approach improves redundancy, load balancing, and operational visibility, and reduces per-node policy maintenance.

[See Non-Router-ID Endpoints in Segment Routing Traffic Engineering.]

SR-MPLS-TE local computation with OSPF underlay (MX204, MX304, MX2010, MX10003, MX10004, MX10008, and MX10016)— We support Segment Routing over MPLS with Traffic Engineering (SR-MPLS-TE) local computation with OSPF as the underlay interior gateway protocol (IGP) for the following features:

- Endpoint IP address configured as anycast IP address
- Segment-list computation using anycast constraints
- Metric-type delay
- Protected and unprotected path selection

[See Segment Routing LSP Configuration.]

SRv6 SR-TE binding SID support (MX204, MX240, MX304, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—The binding segment identifier (BSID) feature segregates the network into multiple traffic engineering tunnels. You can allocate binding SIDs under an SRv6-TE tunnel to reduce the stack size at the ingress node. The ingress node stores only a

subset of the traffic engineering database (TED) from other domains and transit nodes with the binding SID handle convergence. SRv6 binding supports both classic and micro-SIDs (uSIDs).

[See Binding SID (BSID) for Static SRv6 TE Tunnels.]

Support for IS-IS IPv6 SRLG TLV 139 and interface-level SRLG configuration (MX204, MX240, MX304, MX480, MX960, MX2010, MX2020, MX10004, MX10008, and MX10016)—This feature enables routers to advertise and process IPv6 Shared Risk Link Group (SRLG) information as defined in RFC 6119, *IPv6 Traffic Engineering in IS-IS*, ensuring accurate awareness of shared-risk failures. By extending SRLG support to IPv6 and integrating it directly into the IS-IS configuration, you can deploy SRv6 networks without relying on MPLS. This feature enhances Topology-Independent Loop-Free Alternates (TI-LFA) fast reroute path calculation. It enables faster and more reliable protection against link, node, and shared-risk failures for modern IP networks.

To enable per-interface SRLG configuration and advertisement in IS-IS for SRLG-aware TI-LFA, configure `srlg` *`srlg-name`* at the `[edit protocols isis interface interface-name]` hierarchy level.

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

TTL propagation for SRv6 tunnels (MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—We have enabled time-to-live (TTL) or hop limit (HL) propagation for SRv6 tunnels to allow traceroute probe packets generated by customer edge (CE) node to have TTL or HL expiry on SRv6 tunnels. TTL propagation is enabled by default from IP or MPLS to SRv6 (on encapsulation) and SRv6 to IP or MPLS (on de-encapsulation), ensuring traceroute visibility. The TTL or HL is propagated during encapsulation and de-encapsulation between the payload (IPv4, IPv6, or MPLS header) and SRv6 (outer IPv6) header. When a TTL or HL expiration occurs, the transit router generates an ICMP error message and encapsulates it within the SRv6 tunnel. The egress node subsequently de-encapsulates the packet and returns the ICMP error message to the originating source. To disable TTL propagation, use the `no-propagate-ttl` statement at the `[edit routing-options source-packet-routing srv6]` hierarchy level.

To enable tunneling ICMP error messages on transit routers include the `icmp-tunneling` statement at the `[edit routing-options source-packet-routing srv6]` hierarchy level to generate and encapsulate ICMP error packets through the SRv6 path.

[See TTL propagation for SRv6 Tunnels.]

Subscriber Management and Services

External carrier-grade NAT (CGNAT) with RADIUS accounting support for subscribers (MX204, MX240, MX301, MX304, MX480, MX960, MX10004, and MX10008)—You can enable the external carrier-grade NAT feature on a broadband network gateway (BNG) router. With this configuration, you can offload Network Address Translation (NAT) and Network Address Port Translation (NAPT) functions to a dedicated carrier-grade NAT device while retaining subscriber management functions on the BNG router. Use the `external-napt` statement to enable external carrier-grade NAT on the BNG router. You can

also manage external carrier-grade NAT through the RADIUS server by using the predefined variable `$junos-external-napt-control`.

When you enable external carrier-grade NAT and accounting services, the BNG records public IP address and public port block information received from the carrier-grade NAT device through RADIUS vendor-specific attribute (VSA) 238.

[See External CGNAT, cgnat, dynamic-profiles, and Juniper Networks VSAs Supported by the AAA Service Framework.]

PPPoE and DHCP subscriber services support over SoftGRE (MX301, MX304, MX480, MX960, MX10004, and MX10008)—We provide GRE tunnel support for PPPoE and DHCP subscriber services on the following broadband network gateway (BNG) routers:

- MX480, MX960, MX10004, and MX10008 with MPC10, MPC11 or LC9600 line cards.
- MX301 and MX304 with built-in line cards.

Following are the key features:

- Dynamic or static creation and deletion of GRE tunnels
- GRE tunnels on VLAN tagged and untagged interfaces
- IPv4 and IPv6 GRE tunnels on same interface
- VLAN tagged and untagged GRE payloads
- IPv4 and IPv6 DHCP or PPPoE subscriber termination over GRE
- PS over RLT support with both active/active and active/backup modes

[See Wi-Fi Access Gateways, Fixed Wireless Access Solution with Dynamic Soft-GRE Tunnels, and Static Subscribers Over Statically Configured Soft-GRE Tunnels.]

Support for subscriber address preservation (MX204, MX240, MX301, MX304, MX480, MX960, MX10004, and MX10008)—We provide address preservation support for all IPv4 and IPv6 addresses. You can use this feature to save subscriber IP addresses when the subscriber logs out and restore them when the subscriber logs in again. Use the following configuration statements and operational commands to manage this feature:

- `address-preservation` statement to enable preservation of subscriber IP addresses
- `aging-time` statement to configure the duration for holding preserved addresses
- `show network-access address-assignment preserved` command to view the preserved address table
- `clear network-access address-assignment preserved key all` command to delete the preserved addresses

[See Subscriber Address Preservation with Aging Timers, address-preservation (Access), show network-access address-assignment preserved, and clear network-access address-assignment preserved.]

User Access and Authentication

Enhancement: TLS Support for TACACS+ (TACACS+/TCP Secure Transport) (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—This feature enhances the existing TACACS+ over TCP (TACACS+/TCP) implementation by adding support for Transport Layer Security (TLS), enabling secure TACACS+ communication. With this update, the device-side TACACS+ client is extended to establish TLS sessions with TACACS+ servers, ensuring encryption of authentication, authorization, and accounting (AAA) traffic. The enhancement includes TLS handshake handling, server certificate validation, and secure session establishment, while maintaining backward compatibility with non-TLS configurations. This improves overall security by protecting sensitive data from interception and aligns TACACS+ communication with modern security standards.

Sequential DNS fallback for outbound SSH (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When a configured hostname resolves to multiple IP addresses via DNS, Junos OS automatically tries each address in sequence, regardless of IP family, until a connection succeeds or all addresses have been attempted. The `outbound-ssh` retry mechanism activates after all resolved addresses have been tried.

[See outbound-ssh.]

VPNs

Configurable PRF in IKEv2 with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—You can configure Internet Key Exchange version 2 (IKEv2) proposals with a specific pseudorandom function (PRF) algorithm independent of the authentication algorithm. Use the `prf-algorithm` option to configure the following algorithms at the [`edit security ike proposal` *`proposal-name`*] hierarchy level in the IPsec VPN service with the iked process:

- `prf-hmac-sha-256`
- `prf-hmac-sha-384`
- `prf-hmac-sha-512`
- `prf-hmac-sha1`

You cannot combine a PRF with Authenticated Encryption with Associated Data (AEAD) cipher suites such as AES-GCM or ChaCha20-Poly1305 algorithm.

[See prf-algorithm and proposal (Security IKE).]

Deprecation of certificates as authentication method in IKE (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We've deprecated support for the `certificates` option at the [`edit security ike proposal` *`proposal-name`* `authentication-method`] hierarchy level. However, you can continue to use the other authentication methods at this hierarchy level.

[See proposal (Security IKE).]

Multiple local certificates in IKE with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—You can specify multiple local certificates in an IKE policy on a device that runs the iked process for the IPsec VPN service. During IKE negotiation, the device selects the local certificate whose certificate chain matches the peer's certificate request. The device evaluates the certificates in their listed order. Use `set security ike policy` *`policy-name`* `certificate local-certificate [`*`cert1 cert2`*`]` to configure multiple local certificates.

[See certificate and IKE Policy with Local Certificate.]

RSA-PSS signature algorithm in IKEv2 with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Specify the RSA signature scheme for IKEv2 authentication with an RSA local certificate as per RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2*. Configure the following algorithms using the `digital-signature-scheme` option at the [`edit security ike proposal` *`proposal-name`*] hierarchy level:

- `rsassa-pkcs1-v1-5` for RSA signature with the PKCS1 v1.5 scheme
- `rsassa-pss` for RSA signature with Probabilistic Signature Scheme (PSS) as the default signature scheme for RSA certificate

Do not configure `digital-signature-scheme` without setting `digital-signature` as the authentication method.

[See digital-signature-scheme, proposal (Security IKE), and show security ike security-associations.]

Support for trusted CA group for certificate validation in IKE with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Junos OS uses the trusted certification authority (CA) group in an IKE policy to validate peer certificates during IKE authentication. You can use the feature on a device that runs the iked process for the IPsec VPN service. Configure the trusted CA group to specify which CAs the device trusts for incoming peer certificate validation. Use `set security ike policy` *`policy-name`* `certificate trusted-ca trusted-ca-group` *`trusted-ca-group-name`* to reference the trusted CA group in an IKE policy for certificate validation.

[See certificate and IKE Policy with Trusted CA Group.]

Additional Features

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

Chassis management support (MX301)—We support 1+1 redundancy mode for the Power Supply Units (PSUs), Environment Monitoring for sensor components and humidity sensor to display relative humidity and temperature.

[See show chassis routing-engine humidity]

Chassis thermal health check (MX2010 and MX2020)—We monitor the Power Supply Module (PSM) power output and Field Replaceable Unit (FRU) power consumption of your hardware devices and if the PSM power output exceeds the FRU power consumption by a specified threshold, it takes the action you specify (such as auto shutdown or raising alarm). Enable the feature using the CLI command- `set chassis thermal-health-check action-onfail`.

[See thermal-health-check.]

Deprecation of hyper mode (MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—Hyper mode is deprecated. Even if the user configures hyper mode using the CLI, the device will operate in the default normal mode. The following CLI commands previously used to configure the device in normal or hyper mode are retained for software backward compatibility. However, we've marked them as deprecated. These CLI commands display the configured mode and include a message indicating their deprecated status.

1. `set forwarding-options no-hyper-mode`—to operate in normal mode
2. `set forwarding-options hyper-mode`—to operate in Hyper Mode

[See Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches.]

Behavioral changes in firewall filter and policers to support GRE loopback-avoidance mode and GRE-only tunnel services scaling (MX301, MX304, MX480, MX960, MX2010, MX2020, MX10004, and MX10008)—

- Forwarding table filter (FTF) post GRE (Generic Routing Encapsulation) encap/decap will match on older headers.
- Support for firewall filters and policers on GR interface attached using multiple families such as inet, inet6, and bridge.
- Firewall filters attached on an egress logical interface (IFL) will match on new headers post GRE encap/decap irrespective of whether the egress IFL is on the same or different PFE.
- Policers use the packet length corresponding to the current packet for policing after GRE encapsulation/decapsulation. Policers on egress IFLs use the length of the new packet, while older lengths are used for FTF attachment.
- Log/syslog firewall filter actions display corresponding packet snapshots after GRE encapsulation/decapsulation. Log/syslog on egress IFLs displays the new packet, while older packets are shown for FTF attachment.

[See Configuring Forwarding Table Filters.]

Next-hop simplification model in MAC-VRF routing instances with EVPN E-LAN (MX304, MX480, MX960, MX2010, MX2020, MX10004, and MX10008)—We support the next-hop simplification model in MAC-VRF routing instances with EVPN E-LAN.

[See MAC-VRF Routing Instance Type Overview.]

Monitor performance of QSFP-100G-ZR optics (MX240, MX480, and MX960)—The QSFP-100G-ZR transceiver is supported on MX240, MX480, and MX960 with the MPC7E-MRATE line card. Monitor the performance of QSFP-100G coherent ZR optics and receive threshold-crossing alert (TCA) information to efficiently manage the optical transport link. Accumulate performance metrics into 15-minute and 1-day interval bins. Use the `show interfaces transport pm` command to view current and historical performance data.

[See optics-options and show interfaces transport pm.]

Monitor performance of QSFP-100G-ZR-HP optics (MX304)—Monitor the performance of QSFP-100G coherent ZR high power (HP) optics and receive threshold-crossing alert (TCA) information to efficiently manage the optical transport link. Accumulate performance metrics into 15-minute and 1-day interval bins. Use the `show interfaces transport pm` command to view current and historical performance data.

[See optics-options and show interfaces transport pm.]

SNMP MIB walk and traps support for coherent ZR optics (MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—Use SNMP MIB walk and traps to efficiently monitor and manage coherent ZR optics, including 100ZR, 400ZR, 400ZR-M, and 400ZR-M-HP. With this feature, you can retrieve object identifier (OID) information sequentially and receive notifications when alarms are triggered or cleared.

[See Enterprise-Specific MIBs Supported by Junos OS and show snmp mib.]

Support for EVPN-VXLAN Type 5 unequal load balancing (MX304 and MX960)—EVPN-VXLAN Type 5 Unequal Load Balancing uses the EVPN Link Bandwidth Extended Community to advertise Layer 3 link bandwidth and dynamically balance traffic across multihomed devices.

[See EVPN-VXLAN Type 5 Unequal Load Balancing and policy-statement.]

Support for Quantum Buffer in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptographic agility.

[See The Quantum Buffer and moduli.]

Support for Shor-resistant and other default key exchange algorithms in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y,

QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSH supports the Shor's algorithm-resistant hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm. And, by default, SSH supports the `dh-group16-sha512` and `dh-group18-sha512` Diffie-Hellman (DH) group key exchange algorithms.

[See key-exchange.]

Support for symmetric consistent hash (MX301, MX304, MX480, MX960, MX2010, and MX2020)—We now support symmetric consistent hash toward remote dynamic tunnel destinations/MPLS-over-UDP tunnels (MPLSoUDP) across ECMP paths. The feature is available on AFT-based line cards. Note that if you have enabled advanced-forwarding mode, and the device has AFT-based line cards and UKERN-based line cards (which do not support this feature), then UKERN-based line cards will be powered off.

[See Next-Hop-Based Dynamic Tunnels and advanced-forwarding.]

Supported transceivers, optical interfaces, and DAC cables (MX240, MX301, MX480, MX960, SRX4600, SRX5400, SRX5600, and SRX5800)—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.

Two-Way Active Measurement Protocol (TWAMP) monitoring service (RFC 5357) hardware timestamp support to enable flex algo and SR-MPLS support (MX240, MX480, and MX960)—We've extended support to the MPC7E-10G and MPC7E-MRATE line cards. Configure the `offload-type inline-timestamping` option of the `test-session` statement.

[See test-session (Junos OS).]

Monitor network performance and configure user-configurable PM intervals (MX204, MX240, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, and MX10016)—Configure performance monitoring (PM) intervals for long-haul coherent ZR optics across supported platforms to gain more granular and precise visibility. Adjust intervals from the default 15 minutes to shorter durations. Use the `set chassis optics pm interval-length (10s|30s|1min|5min|15min)` command to adjust the intervals according to your specific monitoring needs.

We support the feature on MX304 with MX304-LMIC16; MX2008, MX2010, and MX2020 with MX2K-MPC11E; and MX10004, and MX10008 with MX10K-LC2101, MX10K-LC480, MX10K-LC4800, MX10K-LC4802, and MX10K-LC9600 line cards.

[See User Configurable PM Interval Length, Configure User Configurable PM Interval Length, interval-length (Chassis), optics (Chassis), Coherent Optics Performance Monitoring, and show interfaces transport pm.]

What's Changed

IN THIS SECTION

Learn about what's changed 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 OS Problem Report Search application.

EVPN

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit protocols evpn]` hierarchy level only on platforms that support a default switch instance. To enable

these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the `[edit routing-instances name protocols evpn]` hierarchy level.

[See Preference-Based DF Election.]

- **XML output tags changed for `show l2-learning evpn mcnh-info` command**—We've changed the XML output for this command. The XML `<l2ald-mcnh-entry-list-headers>` element level under the `<l2ald-mcnh-entry>` element has been removed. The child output tags that were under the `<l2ald-mcnh-entry-list-headers>` element are now direct child tags of the `<l2ald-mcnh-entry-list>` element.

General Routing

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the `[edit services port-extender satellite <name> device-model]` hierarchy, you can use the '?' to see a list of supported platforms. We have also added the `skip-lo0-config` statement under `[edit service port-extender jnu]` hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added `cp-mtu` under the `[edit service port-extender satellite <name>]` hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

  [See port-extender (Services).]

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for auto-configuration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address `ff02::1`. The RA source is the router's link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- **Restrict PPPoE client lockout output to underlying logical interfaces with short-cycle protection**—The "show pope lockout" command shows PPPoE clients information only for underlying logical interfaces with `short-cycle-protection` enabled.

  [See show pppoe lockout.]

- **Traffic Engineering Database Import Identifier Update**—The help string for the `set protocols mpls traffic-engineering database import identifier` command has been updated to BGP-TE identifier for standard IGP instance. This change clarifies that the identifier configuration affects only the standard IGP instance. The BGP-LS advertisements in the `lsdist.0` table will now use the configured value instead of defaulting to Instance Identifier 0. This update ensures accurate Instance Identifier

configurations in BGP-LS advertisements, enhancing network visibility and management. The routers running software later than Junos OS release 24.4.x version should verify that the correct Instance Identifier is being advertised. [See import (MPLS Traffic Engineering Database.]

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the `show route extensive expanded-nh` command.

  [See show route extensive <route> expanded-nh.]

- **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.

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)               - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)        -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- **Command to view interface-set information (ACX Series, MX Series)**—Use the `show subscribers interfaces interface-set` command to view interface-set information.

  [See show subscribers interfaces.]

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>>

- **Per-fan entries for fan trays in SNMP MIB tables (PTX10008)**—You can monitor individual fans in a fan tray through distinct entries in SNMP MIB tables. This enhancement addresses the earlier limitation that SNMP MIB does not list individual fans.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.
  - A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Member Link Limit on Adaptive Aggregated Ethernet (PTX Series)**—A maximum of 64 member links are configurable on an aggregated Ethernet (AE) bundle with adaptive load balancing. Configuring more than 64 member links will result in a commit failure.

  [See Understanding Aggregated Ethernet Load Balancing.]

- Write cache disabled alarms no longer raised—The "Disk 2 Write Cache disabled" alarms were reported in the chassis alarms due to a missing SSD model in Junos OS supported device table. The support is now added, and alarms are no longer reported.

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading, copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

  [See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:

  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.

  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.

  - The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

  These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

  [See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Network Management and Monitoring

- **Improved ZTP SNMP configuration handling (Junos)**—We have reordered the SNMP configuration commands to optimize the Zero Touch Provisioning (ZTP) process. The `set snmp engine-id local` statement now processes before `set snmp v3 usm local-engine user`, ensuring SNMPv3 USM authentication and privacy keys generation using the configured local engine ID instead of an auto-derived default. This prevents SNMPv3 query authentication failures from engine-ID resets.

  [See Configure Engine ID on SNMPv3 and Create SNMPv3 Users.]

Platform and Infrastructure

- VXLAN VTEP outer fragmentation in optimized mode (MX Series)—On MX Series routers acting as IPv4-to-VXLAN Virtual Tunnel Endpoints (VTEPs), releases where optimized (loopback-avoidance) mode is the default can encapsulate oversized IPv4 packets into VXLAN and then fragment the outer IPv4 header. This behavior can violate RFC 7348 (VTEPs must not fragment VXLAN packets) and can result in traffic being dropped when the remote VTEP discards fragments. As a workaround, configure loopback mode (loopback-dynamic-tunnel) on the VTEP, increase underlay MTUs to accommodate VXLAN overhead, and ensure PMTUD is enabled end-to-end. Note: If the don?t? fragment (DF) flag is set, drop the packet and send an ICMP ?fragmentation needed? message. If the don't fragment (DF) flag is clear, fragment the inner IPv4 packet before encapsulation.

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **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.]

- **Inbound Convergence Status Added to BGP Neighbor Display (Junos OS and Junos OS Evolved)**—The `show bgp neighbor` command now displays the Inbound Convergence status and cause.

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **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.

User Interface and Configuration

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:

  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`

  - Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`

  - Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

  [See show system data-models.]

- **Introduction of remove operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents <userinput>statement not found</userinput> warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

- **Deprecation of system syslog shell configuration (MX Series)**—We have deprecated the `system syslog shell` configuration hierarchy and scheduled it for removal in a future release. Supported shells log all commands entered into the shell to syslog by default, without requiring any configuration statement.

This shell command logging behaviour now aligns with the CLI, ensuring that all commands are logged.

- **CLI options and RPC formats to retrieve `genstate` information and XPath expressions for operational commands**—A gNMI telemetry collector can subscribe to `genstate` resource paths to query for the corresponding state data on a device. We've added CLI pipe filter functions and RPC formats that enable you to discover available `genstate` resource paths and the associated `genstate` YANG module for a given show command. You can also retrieve the XPath expression for any node in command output. The options include:

  - `| display genstate subscription-paths`—Display all resource paths available for subscription in the `genstate` YANG module associated with the given command.
  - `| display genstate yang-module-name`—Display the name of the `genstate` YANG module that includes the data model and resource paths for the given command.
  - `| display xml genstate-path`—Display the command output in XML and include the `junos:genstate-path` attribute for every emitted tag that is available for subscription. The attribute value is the `genstate` subscription path. For RPCs, use the `format='xml-genstate-path'` attribute in the request tag.
  - `| display xml rpc response-paths`—List all possible response paths under the equivalent RPC's top-level response tag.
  - `| display xml xpath`—Display the command output in XML and include the `junos:data-xpath` attribute for each tag in the output. The attribute value is the XPath expression that selects that specific node. For RPCs, use the `format='xml-xpath'` attribute in the request tag.

  [See | (pipe) and Junos Genstate YANG Data Models.]

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

- **Removal of weak algorithms in IPsec VPN with iked process (MX301, MX304, MX10004, MX10008, MX Series with MX-SPC3, NFX250, NFX350, SRX Series, and vSRX3.0)**—The following algorithms are no longer available as CLI options:

  - `des-cbc` and `3des-cbc` as encryption algorithms in IKE proposal and IPsec proposal at the `[edit security ike proposal encryption-algorithm]` and `[edit security ipsec proposal encryption-algorithm]` hierarchy levels respectively.

- md5 and sha1 as authentication algorithms in IKE proposal at the [edit security ike proposal authentication-algorithm] hierarchy level.
- group1, group2, and group5 as DH groups in IKE proposal and Perfect Forward Secrecy (PFS) groups at the [edit security ike proposal dh-group] and [edit security ipsec policy perfect-forward-secrecy keys] hierarchy levels respectively.
- hmac-md5-96, hmac-sha1-96, and hmac-sha-256-96 as authentication algorithms in IPsec proposal at the [edit security ipsec proposal authentication-algorithm] hierarchy level.
- basic, compatible, and standard as IKE proposal sets and IPsec proposal sets at the [edit security ike policy proposal-set] and [edit security ipsec policy proposal-set] hierarchy levels respectively.

For MNHA HA link encryption and standalone VPNs, using hmac-sha1-96 as the IPsec authentication algorithm at the [edit security ipsec proposal proposal-name authentication-algorithm] hierarchy level results in a commit error. Configure stronger algorithms to enhance IPsec VPN security.

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

Known Limitations

IN THIS SECTION

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

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

General Routing

- In some NAPT44 and NAT64 scenarios, Duplicate SESSION_CLOSE Syslog will be seen. PR1614358

- On unsupported line cards (MPC9 or below), no-tunnel-services knob will get committed and "show interface gr-x/y/z" will show No-tunnel-services knob. However at line card level this will not get applied and PFE will throw following message in syslog regarding the unsupported feature: pfe error log - "no-tunnel-services knob is not supported for ifd:gr-x/y/z on this platform" PR1930878

- 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

- This is seen when LMIC goes for offline/online continuously and results in YT accessing PFE even before PFE is online, This issue will not be seen when PFE is back online, No impact on functional aspects. PR1903848

- When GRE keepalives are enabled on GRE tunnels, packet loss may be observed during a Routing Engine switchover. OAMD process is restarted on the new master RE after GRES, and as per the OAM state machine design, all sessions begin in the DOWN state. The sessions transition to UP only after receiving a GRE keepalive response. Because RPD receives the OAM state as DOWN during this initialization phase, it temporarily retracts the static routes installed over the GRE interfaces. Once OAM transitions to UP, RPD reinstalls the routes. If traffic is using these static routes during the switchover, this withdrawal and reinstallation can result in brief traffic loss. PR1943893

MPLS

- On MX platforms, 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

Routing Protocols

- larger primes of size greater than 8192 are generated when designer type moduli is confiugured with count 3 and 4

  PR1942246

Open Issues

IN THIS SECTION

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

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

EVPN

- On Junos OS, if the router-id is not explicitly configured, the device dynamically selects a router-id during bootup. In this scenario, the device advertises EVPN Type1 AutoDiscovery/Ethernet Segment Identifier (AD/ESI) and Type4 Ethernet Segment Route Target (ES RT) routes with a Route Distinguisher (RD) of 0:0. As a result, Route Targets (RTs) advertised by different Provider Edge (PE) routers end up having identical prefixes. Consequently, only one PE routers RT is advertised and learned. This behavior prevents EVPN all-active redundancy resulting in the loss of multi-homing.

  PR1924472

General Routing

- On MX series platforms, multicast traffic might not be forwarded over User Datagram Protocol (UDP) encapsulation (ud)tunnels if the Automatic Multicast Tunnelling (AMT) is enabled. The multicast traffic might leave the AMT interface, but might be dropped in the Packet Forwarding Engine (PFE).

  PR1638141

- As part of the RSI process, execute the 'show vmhost support-info' command to comprehensively collect the vmhost logs using various cat commands. Some of these commands attempt to access files that do not exist on the MX304, leading to error messages.

  PR1913540

- For ACX5448, MX204 and MX2008 "VM Host-based" platforms, starting with Junos 21.4R1 or later, ssh and root login is required for copying line card image (chspmb.elf for MX2008) from Junos VM to Linux host during installation. The ssh and root login are required during installation. Use "deny-

password" instead of "deny" as default root-login option under ssh config to allow internal trusted communication. Ref https://kb.juniper.net/TSB18224 PR1629943

- 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
- With Next Generation Routing Engine (NG-RE), in some race conditions, the following interrupts messages might be seen on master RE: kernel: interrupt storm detected on "irq11:"; throttling interrupt source PR1386306
- In Netconf private edit configuration session, commit RPC fails when unprotect operation is performed. PR1751574
- Digital Optical Monitoring (DOM) is supported for the optical transceiver modules and not for a copper. However, 'jnxDomAlarmSet' and 'jnxDomAlarmClear' trap will be generated when a copper port becomes down or up.

  PR1560149
- Phase jump of around 300ns is seen upon LAG member switchover from secondary to primary on 400G interface with ZR optics PR1893122
- When a interface supports multiple lanes, the SNMP OID jnxDomCurrentLaneWarnings is incorrectly handled as a single lane, resulting in the value from lane 0 being replicated across all other lanes. PR1906927
- 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
- MX304 is not enforcing the lack of support for stateful firewall PR1924119
- zeroize command not working PR1857029
- JDI-RCT:M/Mx: ISIS session over MPC11 cards flapped due to "3-Way handshake failed" during ISSU (FRU upgrade stage - reboot phase) PR1809351
- Tracing Enabled by Default in JNUD Daemon and the logs are rotated, can be deleted when needed. PR1924811
- On Junos MX Broadband Edge (BBE) platforms where Subscriber access management service Layer-2 Bitstream Access (L2BSA) solution is deployed, when "inner-vlan-id-swap-ranges" is modified in the presence of active L2BSA subscribers, free inner-vlan-id count is not calculated correctly, hence subsequent L2BSA login attempts fail and subscribers move to pending state. PR1941118
- RPC "get-power-usage-information" for cli "show chassis power" shows error wtih ODL response PR1771483

- HTTPv6 redirect functionality doesn't work on the MX301 platform. PR1925736
- BNG user plane: Changing the BNG Controller name is allowed after the user-plane is already associated with the BNG Controller and subscribers are logged in. This leaves the system in a bad state. PR1891388

Interfaces and Chassis

- Before changing the interface-type knob, proactively delete any unsupported IFL configurations based on the new interface-type knob setting. This prevents delay in further commit processing and ensures smooth operation. Please refer the attached document "Understanding GRE IFL Scale Configuration" PR1931166

MPLS

- On MX301, rpd core may be seen after deactivate/activate protocol isis trigger PR1901973

Network Management and Monitoring

- junos/software-information sensor exist in `show agent sensor` output after yang module is deleted

  PR1880338

Platform and Infrastructure

- On Junos OS having the Remote Authentication Dial-In User Service (RADIUS) server with accounting configured, the auditd crash is observed if the Radius server is unreachable for a long time and a large number of accounting records accumulate, which leads to memory exhaustion during accounting processing and results in a process crash. PR1926050

Routing Policy and Firewall Filters

- On Junos OS, rpd process crash is observed and results in core dump when self-referencing routing policy is configured that causes a loop which does not get detected. Traffic impact will be observed.

  PR1943279

Routing Protocols

- In large scale routing instance config, please provide interval of 20-30 mins for system to stabalize PR1883895

Resolved Issues

IN THIS SECTION

Learn about resolved 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 OS Problem Report Search application.

Application Layer Gateways (ALGs)

- On MX platforms with SPC3 line-cards/ SRX platforms, when there are bandwidth attributes carried in Session Description Protocol (SDP) messages inside Session Initiation Protocol(SIP) packets, SIP application traffic and SIP calls get impacted.

  PR1930403

Authentication and Access Control

- On Junos OS platform, management access via telnet disconnects under certain conditions.

  PR1909811

Class of Service (CoS)

- On Junos Fusion platforms, a 54B VXLAN header present on the Next Gen Port Extender (NGPE) extender port for packets going from Aggregation Device(AD) to Satellite Device (SD). These additional packets result in a reduced bps rate than the configured shaping-rate.

  PR1930484

EVPN

- Multiple ifls in sp-style on the same port and part of same bridge domain is an unsupported configuration on evo QFX devices like QFX5130s. This is HW limitation and ASIC (VLAN membership and egress translation) will get mis-programmed because of this configuration. This restriction is applied on both VXLAN and non-VXLAN config. PR1927899

- On Junos OS MX and QFX series platforms, in Virtual Router Redundancy Protocol (VRRP) with Ethernet Virtual Private Network - Virtual Extensible Local Area Network (EVPN-VXLAN) scenario, when VRRP priority is changed, the VRRP virtual Media Access Control (MAC) address can remain pinned as a static entry on the remote device. As a result, MAC movement does not occur correctly, and VRRP packets are dropped on the leaf device. This impacts VRRP operation after priority changes.

  PR1921796

- When Junos and Junos Evolved platform sends two proxied Neighbor Advertisement for default gateway, second one without Router flag and target Link-Layer address. The second Neighbor Advertisement can disrupt connectivity for end hosts. A proxied Neighbor Advertisement (NA) is an IPv6 mechanism where a network device responds to Neighbor Discovery requests on behalf of another device, instead of the actual endpoint replying itself. PR1926468

- On all Junos and Junos Evolved platforms that support EVPN (Ethernet VPN) virtual-switch routing instances, ARP (Address Resolution Protocol) resolution fails for static host routes configured with a /32 mask when the next-hop interface is an IRB (Integrated Routing and Bridging) interface and the destination host resides in a different subnet. As a result, ARP entries are not installed, MAC (Media Access Control) addresses of destination hosts are not learned, and traffic destined to those hosts becomes unreachable, causing service impact. PR1926818

- On all the platforms running Junos and Junos Evolved that support micro-segmentation firewall filter configuration for EVPN GBP, a same firewall filter was not supported by multiple routing-instances.

This put a limitation on using same firewall filter for multiple routing-instances and required users to create same firewall filter with different unique names for each routing-instance. PR1919103

- On all Junos based platforms, in an Ethernet Virtual Private Network- Multi Protocol Label Switch (EVPN-MPLS) setup, when control-word configuration is present in the routing-instance, packets are forwarded on the Integrated Routing and Bridging (IRB) without the control-word flag. This causes the packets to drop at the remote end. There is no impact on data traffic. PR1934165
- In any EVPN setups of MX , EX,QFX using Aggregated Ethernet across multiple FPCs, the VRRP virtual IP may fail to install ARP and destination route entries when the VRRP virtual MAC is learned on one FPC while the ARP response is received on another with a source MAC and SHA mismatch. PR1943641
- router-mac community is not removed when exporting the routes into evpn PR1935768
- On all Junos and Junos Evolved platforms configured with Multiple Protocol Label Switch (MPLS) Virtual Private Network (VPN) services profiles, repeated backtoback routing engine (RE) switchovers can cause an RE0toRE1 switchover to fail to populate Internal Memory Next-Hop (IM NH) table entries, resulting in lookup related errors and impacting the forwarding of Broadcast, Unknownunicast, and Multicast (BUM) traffic and routing engine switchovers which are critical for network stability. PR1910546
- On all Junos OS and Junos OS Evolved platforms running a release from 25.4R1 in an EVPN (Ethernet Virtual Private Network)-MPLS (Multiprotocol Label Switching) or EVPN-VXLAN (Virtual Extensible Local Area Network) environment, with an IRB (Integrated Routing and Bridging) interface configured as the Layer 3 gateway, when ARP (Address Resolution Protocol) probe packets are received while proxy ARP is enabled, the device generates proxy ARP replies for unresolved IP (Internet Protocol) addresses. This causes false duplicate IP detection and leads to DHCP (Dynamic Host Configuration Protocol) address assignment failure. PR1947816

Forwarding and Sampling

- On specific VMHost platforms like MX204/MX304/MX10003/MX10008/MX10016/ PTX1000/ PTX10008/PTX10016/ EX2300/EX3400/EX4650/ QFX5100/QFX5110/QFX5200/ QFX10002/ QFX10008/QFX10016/ vSRX/SRX1500/SRX1600/SRX2300/SRX4100/SRX4200/SRX4300/ SRX4600 platforms. Whenever a client (e.g., lacp/ifinfo/snmp) requests statistics from PFE, the query is routed via the kernel and not through the use of the BULKGET protocol. PR1914977
- When the CLI setting "set forwarding-options family inet route-accounting" is enabled to activate IPv4 route accounting statistics, the streaming for /interfaces/interface/subinterfaces/subinterface incorrectly reports all IPv4 counters as zero. This occurs even though IPv4 traffic is successfully forwarded and accounted for in the Packet Forwarding Engine (PFE) . PR1917717
- On MX304 platforms, when an LMIC (Line-Card Module Interface Card) restart/swap occurs and firewall filters with interface or ARP (Address Resolution Protocol) policers are applied across interfaces belonging to two different LMICs, a policer address collision across PFE (Packet

Forwarding Engine) partitions occurs, leading to discard filter installation and resulting in traffic drops until the FPC (Flexible PIC Concentrator) is restarted to restore normal forwarding. PR1951120

General Routing

- On all Junos platform in MH (Multi homed) EVPN (Ethernet Virtual Private Network) MPLS (Multi Protocol Label Switching) topology where Egress Link Protection feature is enabled on the PEs (Provider Edge) routers and the PE - CE (Customer Edge) link is disabled on the one of the MH PEs and 'restart l2-learning' command is executed on another MH PE, this results in broadcast and Multicast traffic drop. PR1912050

- On MX platforms with MPC11, LC9600, LC4800 Linecards, MX304 and all Junos OS Evolved platforms, optics modules intermittently reports extremely high false temperature values in "show interfaces diagnostics optics" command. This can trigger false high-temperature alarms, leading to unnecessary optics shutdown. PR1937870

- Drop of traffic to subscriber DHCPv6 prefixes may be observed on LNS (L2TP network server) if CPE uses IPv6 address obtained via NDRA process as the source address for DHCPv6 negotiation instead of link-local address. PR1922536

- In Junos, MX10003 platforms with Federal Information Processing Standards (FIPS) enabled experience repeated halts due to Known Answer Test (KATs) failures in the SMIC_QSFP28_MACSEC_TIC crypto path. PR1905490

- On all Junos and Junos Evolved platforms supporting 'chained-composite-next-hop' for BGP-LU (Border Gateway Protocol - Labeled Unicast), when having L2Circuit (Layer 2 Circuit) links configured, with indirect next-hop resolving on BGP-LU prefix, if 'chained-composite-next-hop' enabled for BGP-LU and with 'preserve-nexthop-hierarchy' configured, L2Circuit traffic will get dropped due to failure in installing the forwarding next-hop with labels for L2Circuit service. PR1931875

- FPC crash and aftd-trio core-dump will be observed on MX platforms supporting MPC10E, MPC11E, LC9600 line cards, and MX304 after configuration changes were made to a service-filter which was in use by BBE subscribers. PR1928462

- On MX 10000 platforms, when PTP is configured with a large number of master streams (128 or 256), the time error performance degrades. PR1908234

- On Junos platforms, a race condition occurs when the Simple Network Management Protocol (SNMP) poller attempts to retrieve device license information while subscriber management and SNMP policing features are in use. This condition causes the license-check process to encounter a null pointer and generate a core dump. There is no traffic or service impact. However, license information retrieval through SNMP polling fail until the license-check service recovers. PR1925706

- A cli-pfe show command may continue to gather data in the background after the user issues a Ctrl-C. Eventually the background command will complete. However, the CPU usage for the cli-pfe process will continue to be high while it is still running. PR1915225

- In an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) deployment, a remote MAC move event can cause the remote MAC IP entry to be removed after it is initially installed, resulting in missing Address Resolution Protocol (ARP) state and incomplete hardware programming on the remote Virtual Tunnel Endpoint (VTEP), which can lead to traffic loss for the affected endpoint. PR1936648

- The issue is seen on all MX platforms with MPC10/MPC11/LC4800/LC9600 linecards and MX304, when the output filter and interface policer is applied to the egress AE (aggregated ethernet) and physical interface in the output direction and the interface policer runs first then the output filter. This kind of an incorrect policing issue leads to the exhaustion of the policer bandwidth and throughput reduction will be seen. PR1934419

- On JunOS MX240/MX480/MX960/MX2008/MX2010/MX2020 platforms with MPC10E/MPC11E line cards as well as the MX304 platform, an issue in IPv6 EVPN (Ethernet Virtual Private Network) during NDP (Neighbor Discovery Protocol) resolicitation of a link local target address causes the system to use the IRB link local address in outgoing packets. Stack corruption happens when handling the IRB link local address and resulting in continuous FPC (Flexible PIC Concentrator) reboots (around 3 or 4 times) and crash files generation. PR1928530

- On Junos MX301 and SRX4700 platforms, there is the change in log message for the CB 0_TEMP_0 sensor. This is a display issue with no functional impact. PR1924853

- On Junos OS, MX Series (MX240, MX480, MX960) devices that use a Switch Control Board (SCB), forwarding traffic may be interrupted for approximately one minute if an SCB is physically removed without first powering it off. The control plane remains operational, but forwarding is impacted until fabric manager and forwarding services selfrecover. PR1911536

- l2ald Core @l2ald_gbptag_change_mac_in_db, @gbptagdb_update_internal_taginfo is seen in Suite teardown PR1925496

- 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

- In fusion setup on MX platforms, in a NGPE(next-gen port-extender) topology, If the MTU on an NGPE EP is set to more than 1522, and there are host bound packets of size are than 1522, this will cause the PFE crash. This will cause the system to go down as the connectivity to SD will be lost. PR1922669

- License Feature used is shown 0 for MACSec Bandwidth Usage (in gbps), although MACSec license is installed and used. PR1921529

- In MX uKern and AFT based cards, NH IFL throttling will be seen when resolution for an destination IPv4/IPv6 address fails. PR1927194
- On Junos platforms with MPC 1-9, LC2101, LC480, LC2103 line-cards and MX204, a slow memory is observed when subscribing to the pipeline sensor (path is /components/component/integrated-circuit/pipeline-counters/). Over time, this causes a gradual increase in heap memory usage. Increase in heap memory usage beyond a certain threshold will lead to FPC crash resulting in loss of services and traffic over that FPC. PR1921361
- On Junos MX2020 platforms with SFB3, when event script is configured on the device, the platform does not correctly apply the SFB power-redundant-mode configs (set chassis power sfb-redundant-mode) during a full chassis reboot, affecting the chassis resiliency. PR1924872
- On all Junos OS platforms with dual Routing Engine, an incorrect IPv6 next-hop configuration in a static MPLS LSP triggers a ksyncd (Kernel Synchronization Daemon) replication error on the backup Routing Engine and leads to live kernel core and a ksyncd core files without any service impact. PR1920718
- On Junos OS and Junos OS Evolved platforms that use transport-class and Class-Based Forwarding (CBF), when a service route is installed in an any-class routing table, the service route resolves over a Label Information Structure Table (LIST) next hop (multipath). In this scenario, the total summary count is incremented, but the UP count is not incremented. To maintain consistency with other routing tables, the total tunnel count increment is also blocked. As a result, the total count is incremented in the show dynamic-tunnel summary output for tunnels in the junos.any table, while the corresponding UP count is not incremented. PR1917815
- On MX301 platform, After SFP (Small Form-factor Pluggable transceiver) Online Insertion and Removal (OIR) on a MACsec (Media Access Control Security) 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
- On Junos MX301 and MX304 platforms there is a scaling issue to configure more than 1024 ifls (logical interfaces) on 1-Gigabit ethernet interfaces. Further interfaces will not work. There is no impact on the existing interfaces and system works fine with 1024 logical interfaces. PR1921080
- PR1924183
- On all Junos and Junos OS Evolved platforms, when FIPS is enabled, the JSD server fails to start if a local SSL certificate is configured on telemetry; as a result, telemetry services using SSL certificates are unavailable. PR1902691
- The JavaScript Object Notation (JSON) encoding of leafs of type "ieeefloat32"(https://github.com/openconfig/public/blob/master/release/mode ls/types/openconfig-types.yang#L127) is not correct, causing gRPC Network Management Interface (gNMI) data outputs unreadable. PR1923848

- On Junos OS Evolved ACX and PTX platforms with Dual RE, a configuration commit will not sync to the standby RE (Routing Engine). If a switchover happens during this time, traffic will be impacted because both REs are not in sync. This is a rare case. PR1890518

- In Show dhcp-security binding detail command output, IP vlan and Interfaces are printed first. Older version looked like this: Client IP Address: 40.1.10.1 Logical System: default Routing Instance: default-switch MAC Address: 00:10:94:00:00:01 Vlan: v100 Lease Expires in: 421114 State: BOUND Interface: ge-0/0/34.0 New output looks like this: Client IP Address: 40.1.10.1 Vlan: vlan100 Interface: mge-0/0/7.0 Logical System: default Routing Instance: default-switch MAC Address: 00:10:94:00:00:01 Lease Expires in: 421200 State: BOUND PR1937652

- On all Junos and Junos OS Evolved platforms, after a system reboot, sending a very large number of RPC (Remote Procedure Call) calls per second (20K or more) will prevent gNMI (gRPC Network Management Interface), gNOI (gRPC Network Operations Interface) and P4RT connections from being established. This will result in nginx session failures. There will be no traffic impact due to this issue. PR1915308

- During update and downgrade of vmhost images the livirtd.conf cleanup entries had not been performed properly. See TSB103739. https://supportportal.juniper.net/s/article/VM-Host-system-will-fail-to-sta rt-after-a-system-reboot-due-to-expired-Vcertificate PR1924890

- On all Junos platforms, IPv6 packets with Hop-by-Hop extension header(3<length) with FrameSize 1471 bytes or longer are dropped in Routing Engine. PR1931587

- When AE interface (member links) are configured with MACsec, the MKA session establishes successfully when the default standard MAC address is used. However, when a custom unicast MAC address is configured, the MKA session does not come up. This is not a common use case, but the workaround is to allow the MKA session to operate with the default MAC address instead of a custom unicast MAC address. PR1909930

- Dual stack PPPoE/DHCPv6 syslog reports PFE_ERROR_INVALID_STATE fd01:3333:b205:44f1::/64 => source lookup is not enabled for flow 18558885 PR1924812

- On all Junos platforms that use the NodeSlice architecture like MPC10 and MPC11 FPCs, the issue involves traffic loss, specifically for traffic flowing from R2(Router 2) to R1 (Router 1), while traffic from R1 to R2 flows without issues. Due to this multicast traffic over EVPN(Ethernet Virtual Private Network) MPLS(Multiprotocol Label Switching) maybe blackholed. PR1922876

- On MX10k platforms with LC480 , LC4802, MX304, traffic loss is seen on all ports after disabling and enabling PTP (Precision Time Protocol) configuration. The services will self-recover when traffic recovers in few milliseconds.

  PR1904131

- On MX304 platforms/LC4800/LC4802 linecards, when LMIC(Line Modular Interface Card)/PIC (Peripheral Interface Card) is restarted on MACsec (Media Access Control Security) enabled interfaces, LMIC/PIC will be stuck in offline state which results into traffic loss. PR1922064
- Bad checksum packet is not counted for L3 incomplete PR1929899
- On Junos platforms with Micro-BFD and LACP (Link Aggregation Control Protocol) configured, if Micro-BFD (Bidirectional Forwarding Detection) is deactivated on the peer device and the member interfaces on that peer are flapped, the AE (Aggregated Ethernet) interface remains down even after the member links come back up. Due to this, the AE will be down and traffic loss occurs temporarily. PR1912709
- On Junos MX304 and MX10K platforms, the 'terminal' option is not supported in controller satellite setup along with device-list option since terminal is local to controller and cannot be used to install license in satellite. The user can use the 'file' option and the file needs to be copied to the device and later user can use the same command with the file option to add the license. PR1932385
- On Junos MX platforms with MPC1-9/LC480/LC2101 (including MX204 and MX10003 platforms) when "set protocols l2circuit resolution preserve-nexthop-hierarchy" is enabled, FPC heap memory will leak when CCNHs (Chained Composite Next Hops) are recreated because of the VPLS (Virtual Private LAN Service) PNH (Preserve Nexthop Hierarchy) are getting deleted and re-added. PR1922741
- On MX10004 and MX10008 platforms where SNMP (Simple Network Management Protocol) is configured, jnxLEDState mib displays default value other(1) for PEM LED when "show snmp mib walk jnxLEDState" is executed. This is a display issue and there is no service impact. PR1922744
- On Junos OS MX240/MX480/MX960 platforms with MX-SPC3 cards, the flowd process crash and reboots the service PIC when a stale IPv4 session is mistakenly reused for IPv6 due to a memory issue. During process restart the services remain down, session information is briefly lost and active traffic will drop, however the system recovers automatically. PR1925039
- On Junos platforms supporting Ethernet Segment Identifier (ESI) based Aggregated Ethernet (AE) Link Aggregation Group (LAG) deployments, Link Aggregation Control Protocol (LACP) does not return to the expected Attached state when a Graceful Routing Engine Switchover (GRES) occurs on the Provider Edge (PE) device and is followed by an AE member interface flap on the connected device, with lacp-oos-on-ndf (LACP out-of-sync on non-designated forwarder) enabled on the PE. PR1938565
- On MX240, MX480, MX960, MX2008, MX2010, MX2020 platforms with MPC2/3/4/5/6/7/8/9 line cards, the new subscriber login request will fail if subscriber login request is initiated when PFE (Packet Forwarding Engine) is booting. PR1942159
- On certain MX Series and SRX Series platforms running Junos, an intermittent PLL (Phase-Locked Loop) lock loss on certain line cards leads to fabric link errors. When the issue happens, forwarding

plane traffic across the chassis fabric is impacted, resulting in partial or complete traffic loss. PR1948588

- On Junos OS and Junos Evolved OS platforms, with on_change streaming telemetry subscription mode enabled, if a change is made on a BGP (Border Gateway Protocol) external community (ext-comm), the telemetry on_change module does not generate an update. PR1934534
- On MX platforms with LC9600, LC4800, LC4802 and JNP304-LMIC Line cards, source MAC (Media Access Control) learning may fail on interfaces associated with remote odd PFE (Packet Forwarding Engine) instances when the corresponding even PFE instance is not part of the Layer 2 domain. As a result, traffic destined for these MACs may be flooded instead of being forwarded, even though the source MAC is learned on local PFE. In large scale network, this may lead to performance issues. PR1930172
- On certain MX platform, there is an issue in scaling PPPoE sessions from approximately 131K to 384K when operating in an SVLAN subscriber model. PR1936257
- PFE error messages are observed on MX304 and EX92K when certain kind of L3 unicast and L3 multicast traffic are transmitted PR1907722
- A software issue was identified in the license check process that could cause the process to crash due to a synchronization condition between internal threads. This behavior may occur when multiple threads access the LicenseMonitorEventNotify() function simultaneously, potentially leading to a mutex deadlock and generating a core dump. This issue has no impact on traffic forwarding or network services. PR1933553
- On all Junos platforms, if an interface is configured with "input-vlan-map tag-protocol-id" and the associated "ethernet-switch-profile tag-protocol-id" is deleted while the VLAN mapping is still active, the system may incorrectly allow the configuration commit. If the device is subsequently rebooted or the chassis is restarted, the device may fail to recover and remain down, resulting in traffic impact. PR1933501
- Currently, service-set names can be configured without a strict character limit. However, when 48 characters or more are used, an internal fixed-length limitation causes a buffer overflow. For example, if the same rules are applied to both input and output, the configuration may no longer remain uniquely identifiable. In such cases, the router may interpret the configuration as duplicated, resulting in a crash and generation of a core dump. * If the same service-set with a name of 48 characters or more is configured for both input and output, a core dump will always occur. * Note that using 47 characters does not guarantee that this issue will never occur. Since a character overflow has already happened internally, a core dump may occur even with fewer than 48 characters depending on conditions. * Because variable-length identifiers are also involved, it is safer to keep the name within 40 characters or less to prevent overflow. After the fix in this PR, service-sets exceeding 42 characters can no longer be configured. In addition, the fix also takes identifiers into account, so overflow is highly unlikely even at 42 characters. PR1951987

- On MX304 platforms, PEM Logic Fault alarms may intermittently appear due to I2C transaction timing limitations. PEMs require a delay between backtoback transactions; without this delay, alarms can be raised randomly on either PEM. PR1923135
- On MX platforms with LC2101, LC2103, LC480, MPC1, MPC1E, MPC2, MPC2E, MPC3E, MPC4E, MPC5E, MPC6E, MPC7E, MPC8E, and MPC9E line cards, a Packet Forwarding Engine (PFE) restart is observed when sensor-based statistics are enabled (for example, using set protocols mpls sensor-based-stats) and a rare internal timing condition is encountered during sensor processing. PR1924105
- In certain Junos OS releases, BGP sessions may repeatedly reset when IPv6 Flow Specification (inet6-flow) is enabled. This behavior occurs when the router receives specific IPv6 Flow Specification updates that include a default (wildcard) match. Earlier releases do not exhibit this behavior, leading to a difference in observed stability when upgrading. PR1945627
- On all Junos and Junos Evolved platforms, an enhancement to remove the commit check for uplink protection for the L3 tracking interface. PR1895408
- On all MX10004/MX10008/MX304/MX301 platforms, the three packet types under "protocol exceptions" in the DDoS flow detection configuration?("unclassified," "mtu-exceeded," and "mcast-rpf-err") are hidden commands. Even if these commands are configured, the system reports them as "unsupported platform."If the global flow-detection mode is set to "on" or "automatic," the feature does operate, so flow detection itself is available.However, since the individual settings are not applied, it is not possible to enable flow detection for only specific protocol exception packet types. PR1927093
- The VSTP (VLAN Spanning Tree Protocol) bridge-priority configuration does not apply to VLAN (Virtual LAN) groups as expected. This issue results in VLAN groups not reflecting the intended bridge priority. This configuration issue can affect the network's root bridge election process, causing potential inconsistencies in network behavior and performance. PR1938991
- This is not a major change which will impact the functionality of the system. It is minor change to fix a CLI configuration change. PR1869720
- When using an SRv6-TE path with seven or fewer hops, packets exported from the ingress PE contain extra empty containers within the Segment Routing Header (SRH). This issue does not occur when the SRv6-TE path consists of eight or more hops. PR1947531
- On all Junos and Junos Evolved platforms with PCEP (Path Computation Element Protocol) configured, the pccd process crash is seen when multiple LSP (Label Switched Path) updates are received in a single PCE (Path Computation Element) message. The LSP update sent by the PCE will not get programmed into the forwarding plane and LSP states will not be updated till the pccd process restarts. PR1929225
- On Junos and Junos OS Evolved platforms where multicast forwarding is enabled, a Routing Engine (RE) switchover, whether planned or unplanned, results in temporary multicast packet loss when

active multicast streams are present. This issue is observed only when Protocol Independent Multicast (PIM) is configured together with Internet Group Management Protocol (IGMP) for IPv4 or Multicast Listener Discovery (MLD) for IPv6. During the RE switchover, multicast forwarding state is temporarily removed from the forwarding plane and subsequently restored. While the forwarding state is being rebuilt, multicast packets are dropped for the affected multicast groups. Unicast forwarding is not impacted, and multicast control-plane sessions remain established throughout the event. PR1929521

- On all Junos VMhost platforms, when Junos performs disk access, a rare deadlock involving FreeBSD kernel processes will occur. This condition causes FreeBSD kernel panic, resulting in crash of vmcore and an unexpected device restart, leading to temporary system unavailability. PR1927342
- On all Junos platforms with SSD (Solid-State Drive) of Innodisk and model no. is DGM28-B56M71EC1QF-B5531, alarms are observed for secondary SSD along with continuous input/output ata timeout errors for secondary disk read accesses when the SSD firmware is unable to can't handle the continuous metadata reads. No service impact will be seen as the alarms are seen for secondary SSD alone. PR1933347
- smid daemon may crash with core dump during long-running sessions. PR1944925
- On all Junos MX platform, when MPC10E interfaces use breakout cabling, administratively disabling one breakout port may unexpectedly impact another port within the same MAC quad due to nondeterministic breakout initialization. PR1945343
- On all Junos OS Evolved platforms with coherent optics (ZR/ZR-M/ZR-M-HP), when interface speed configuration is changed along with the optics application parameters in quick succession before the optics get initialized, the picd (port interface concentrator daemon) crashes, resulting in the interfaces going down. PR1931835
- On all Junos OS platforms, configured with VPLS (Virtual Private LAN Service) and IRB (Integrated Routing and Bridging) interfaces, when the MTU (Maximum Transmission Unit) configured on the IRB or CE interface is larger than the MPLS (Multiprotocol Label Switching) core interface MTU (including MPLS label overhead), ICMP Echo Request (ping) packets with higher packet sizes fail. PR1933452
- On all Junos and Junos OS Evolved platforms with Border Gateway Protocol (BGP) signalled Segment Routing (SRv6) dynamic tunnel and connected to a third party vendor device, when an update packet with SRv6 Service Sub-TLV (SSTLV) length as zero is received from a peer, BGP SRv6 dynamic tunnel resolution fails. As tunnel resolution fails, traffic intended for this SRv6 path will be dropped. PR1926805
- On Junos platform in an extremely rare circumstances due to internal race conditions the jinsightd process crash is observed. J-insight Monitor will not function until the jinsightd process recovers. However, there will be no traffic impact due to this issue. PR1935089

- On EX4400 platforms, loss of PoE (Power over Ethernet) power is seen on ports after port bounce or even during normal operation without specific external trigger . This causes outage on all PoE ports and devices connected to PoE port will not receive power causing service impact. PR1930080
- On platforms supporting dot1x authentication, when a Radius accounting Interim-Update is manually triggered from the switch (using a COA), the resulting accounting packet does not include the Framed-IP-Address attribute. However, when the same session generates an interim update via the periodic timer (10-minute interval), the Framed-IP-Address is included as expected. PR1934680
- PR1908428
- On Junos MX platforms with line card scale, and subscriber over AE (Aggregated Ethernet) is enabled, if one FPC hosting an AE member link restarts, another FPC hosting a different member link of the same AE bundle will crash. When the crash occur, PFEMAN is stuck for more than 240 seconds and traffic loss will occur on the affected FPCs. PR1929254
- On all MX platforms where Enhanced Subscriber Management (bbesmgd) is enabled, activating PPPoE (PointtoPoint Protocol over Ethernet) over L2TP (Layer 2 Tunneling Protocol) triggers repeated crashes of the bbesmgd process, resulting in core dump generation. This issue affects both Junos and Junos EVO platforms. PR1939233
- On Junos OS and Junos OS Evolved platforms, an issue has been identified where post-dated licenses do not automatically activate when their configured start date is reached. This issue occurs when a new license is installed with a future activation date. Once the start date arrives, the license remains inactive instead of being automatically enabled as expected. PR1950936
- On Junos OS MX Series platforms supporting Junos node slicing, when multiple FPCs are repeatedly deleted and reassigned across the same or different GNFs, an internal race condition will prevent system events related to FPC deletion and reassignment operations from being processed. As a result, PICs on the affected FPCs fail to come online and remain down, impacting traffic forwarding on interfaces hosted on those PICs. PR1952112
- On all MX platforms with MS-MPC/MPC-3D and MPC1-MPC13E line cards, when Class of Service (CoS) is configured on an Aggregated Ethernet (AE) interface, a crash is observed in the smpc_re process during PPPoE subscriber login attempts. Subscriber management services will be impacted. PR1933096
- On all Junos VMHost platforms, stale SSH login sessions can cause commit and commit-check operations to intermittently stall at the DCD(Collecting status of the Interface Control process) stage. PR1927165
- MX301:: Remove num-of-subports 1 config on SFP ports for 10G and 25G speeds PR1930331

Infrastructure

- vmhost upgrade may fail due to 'Validation failed' when upgrading to freebsd15 based Junos from older releases. This does not apply to Junos-EVO. PR1932893
- On all Junos OS Evolved platforms, due to linux kernel issue, periodical sync from system time to RTC (Real-Time Communication) every 11min stops working when system time moves to backwards. PR1939674

Interfaces and Chassis

- On Junos and Junos OS Evolved platforms with Dual RE (Routing Engine) systems or VC (Virtual Chassis) setup, in an MPLS (Multi-Protocol Label Switch) network along with LDP (Label Distribution Protocol) signalling protocol and P2MP (Point to Multipoint) is used and recursive FEC (Forwarding Equivalence class) is configured such as 'set protocol ldp p2mp recursive fec' on the LER (Label Edge Router)/PE (Provider Edge) device, the multicast traffic statistics doesn't show as expected. Additionally, in the above setup, if there is RE switchover performed or happens due to reason such as component failure, kernel issue on primary RE or the primary RE fails (but not limited to) on the LER / PE the multicast traffic drop is observed. However, the multicast traffic drop happens only during the brief moment of RE switchover. Once the switchover is completed and the label is reprogrammed, the traffic recovers automatically, resulting in minimal traffic loss. PR1917186
- On wan-phy interfaces, show interfaces CLI takes long time (~10 seconds) and SNMP get on interface-mib(IF-MIB) times out PR1941128
- On all Junos and Junos Evolved platforms, after vrrpd restart one out of 2000 vrrp session got stuck in bring up state because the link local address update which is expected after vrrpd restart had failed for that IFL (Logical interface). PR1942571

J-Web

- When a MAC-RADIUS authentication succeeds and the RADIUS server returns CWA redirect attributes (URL-Redirect), the EX switch should intercept client HTTP traffic and return an HTTP 302 redirect to the Captive Portal. The switch instead responded with HTTP 405 (Method Not Allowed), causing the client to bypass the CWA redirect workflow. PR1931780

Layer 2 Ethernet Services

- On all Junos and Junos Evolved platforms, when a loop occurs in the transmission switch, the device starts receiving looped LACP (Link Aggregation Control Protocol) PDU's from itself, instead of messages from the actual peer device. This causes the system to mistakenly believe that a valid LACP connection exists, even though the peer device is not actually connected.As a result, it continues to forward traffic as if the peer were active. Since no valid peer connection is present, this can lead to traffic blackholing . PR1874126

- On all MX platforms, when configuring Dual Stack DHCPv6 ALQ without configuring DHCPv4 ALQ (active-leasequery), the ALQ will not be able to establish the TCP (Transmission Control Protocol) session for the ALQ peer due to XID (Exchange Identification) mismatch. This can lead to DHCPv6 subscriber bindings sync issue and server ALQ, resulting in compromising the redundancy functionality. PR1934742

- jdhcpd cores seen after upgrade with dhcpv6 for IANA/PD PR1927449

- Insert command not working for "dhcp-local-server group" hierarchy, below error is seen: user@host# insert system services dhcp-local-server group servers-1 after syntax error, expecting `after' or `before'. user@host# insert system services dhcp-local-server group servers-1 bef syntax error, expecting `after' or `before'. PR1939685

- Additional syslog messages seen when connected to non-Juniper peers that continuously change the vendor-specific fields in LACP PDU. PR1926024

Layer 2 Features

- An rpd core triggered on the backup Routing Engine after a switchover due to an assertion failure in within the context of L2VPN VPLS auto-site processing, caused by a mismatch between the auto-site claim-id/site-id and the local site-id associated with interfaces in the VPLS IFL repository for a given instance, leading to stale interface entries being referenced during auto-site processing. PR1885690

MPLS

- On Junos OS and Junos OS Evolved platforms with Graceful Restart and RSVP-TE (Resource Reservation Protocol - Traffic Engineering) configured, an rpd crash is observed, leading to traffic impact after a Graceful Restart if a PVC (Permanent Virtual Circuit) fails to allocate correctly during this recovery process, leaving it in an incomplete or unallocated state, and the system attempts to clean up or remove this unallocated PVC. PR1923867

- On all Junos and Junos OS Evolved platforms, MPLS (Multiprotocol Label Switching) traceroute shows noncompliant for routes in a nondefault routing instance because it looks in the wrong routing table, even though the network path is working correctly. PR1943768

- On Junos and Junos OS Evolved platforms, the rpd process crashes when Resource Reservation Protocol (RSVP) sessions are repeatedly cleared under specific timing conditions. This issue occurs due to a timing window during RSVP processing on the standby Routing Engine (RE) while sessions are being re-established, which results in an invalid memory access and leads to an the rpd process crashes. PR1913565

Network Management and Monitoring

- On all Junos and Junos Evolved platforms, with an established telemetry session with multiple collectors, there will be interface statistics discrepancy when there is a linecard OIR event or PFE offline/online sequence. PR1913260

- In all Junos and EVO platforms, a gRPC connection between the SWAN (Software Wide Area Network) agent and the device is being closed unexpectedly. While SWAN client/JET(Juniper Extension Toolkit) client is connecting and sending data close to 3K(3000) routes at 8KB Per RPC (Remote Procedure Call) over the connection, they can hit rare race condition which can cause the clients to disconnect. PR1919448

- System connections are not established while Subscribing for OC path using ssl local-certificates after upgrading to 24.4R2-S3-EVO. This issue is only applicable on 24.4R2-S3-EVO and 24.4R2-S4-EVO PR1933395

Platform and Infrastructure

- On Junos and Junos OS Evolved platforms, the output format of the fetch command changed after upgrading to a newer upstream version of the utility. The earlier format displayed an explicit completion indicator (for example, 100%). The upgraded version removed this indicator. PR1924159

- On MX platforms , multicast traffic over VLAN interfaces (IRB) will be impacted after certain CoS (Class Of Service) configuration changes. This can result in multicast traffic not reaching downstream receivers. PR1950853

- On MX platforms using line cards such as MPC 1-9, LC2101, LC480, LC2103 and MX204 when a Traffic Control Profile (TCP) is configured at both the Interface Logical Set (IFLSET) and the Interface Logical Interface (IFL) levels without an explicit shaping rate, the scheduler could incorrectly apply a shaping rate of zero to the IFL instead of inheriting the shaping rate from the associated Interface Physical Device (IFD). As a result, traffic destined for the affected IFL may experience unexpected packet drops due to the zero shaping rate applied at the scheduler level. PR1932824

Routing Protocols

- On all Junos Platforms, a rare memory corruption condition may occur in the Packet Forwarding Engine (PFE) when Spanning Tree Protocol (STP) is enabled and operating in default (distributed) mode. When the issue is triggered, the PFE crashes and a dc-pfe core file is generated. The exact trigger for the memory corruption is currently unknown. PR1931633

- 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

- NGMVPN multicast traffic is dropped post changing the route distinguisher and route target on L3VPN instances. This issue is intermittent and occurs when RD and RT are changed in all routing instances while traffic is active. In some cases, the issue is observed only in specific instances. When the issue occurs, traffic drops for the inclusive tunnel, affecting all flows on the inclusive tunnel. Recovery may not occur unless external events, such as a route change, take place. PR1909265

- AS path loop may not be detected with VPN route when autonomous-system independent-domain is configured with 'no-attrset' option under routing-options hierarchy PR1926835

- On Junos OS, when Long-Lived Graceful Restart (LLGR) is configured and a Border Gateway Protocol (BGP) neighbor goes down and reconnects after the LLGR stale timer expires, the Routing Protocol Daemon (rpd) process crashes leading to service disruption.

  PR1915893

- On all Junos platforms when GRES is enabled and any other protocol dependent on Periodic Packet Management (PPM) in non-centralised mode(distributed to FPC's) is configured, after an RE (Routing Engine) switchover, the protocol fails to recover, resulting in services dependent on the protocol to go down.

  PR1935823

- On all Junos OS platforms, telemetry data collection after an interface flap event, on an interface with PIM (Protocol Independent Multicast) sparse-dense mode configured causes rpd (Routing Protocol Daemon) process crash and core-dump generation. This results in protocol restarts and service disruption.

  PR1937530

- On the Junos MX2008 VMhost platform with dual Routing Engines (REs), the backup SPMB {Switch Processor Mezzanine Board} (spmb1) fails to come online following a graceful switchover on releases starting when upgrading to releases starting from Junos 24.1. As a result, if the backup SPMB cannot boot, the traffic will be impacted during switchovers. PR1922002

- On all Junos OS platforms with dual RE & Nonstop Active Routing (NSR) enabled, continuous JTASK_IO_CONNECT_FAILED errors on backup RE will be seen. This will not have any functional impact.

  PR1939596

User Interface and Configuration

- When using the load override command, committing a scaled configuration on top of an already committed scaled configuration may take longer to complete. During this time, other sessions cannot enter configure mode, causing those sessions to appear unresponsive. Use the load update command instead of the load override command to load the configuration before committing. PR1906880

- The [system syslog shell] configuration hierarchy is now deprecated and will be removed. All shells that Junos ships will always have syslog commands entered into them as baseline functionality. PR1947442

VPNs

- On Junos OS platforms, traffic drop disruption can occur in multi-homed Next Generation Multicast Virtual Private Network (NGEN MVPN) setups when a Provider Edge (PE) acting as both multicast

source and Designated Router/Rendezvous Point (DR/RP) experiences path changes or flaps, provided the source is local and multi-homed across two PEs.

PR1918004

- On all Junos platforms that support both InService Software Upgrade (ISSU) and NextGeneration Multicast VPN (NGMVPN), multicast traffic loss can occur during an ISSU operation when NGMVPN deployments use both the Inclusive Provider Multicast Service Interface (IPMSI) and the Selective Provider Multicast Service Interface (SPMSI) with a nonzero threshold. During the ISSU process, Flexible PIC Concentrators (FPCs) upgrade and synchronize, and multicast statistics may briefly report zero. This event causes withdrawal of the SPMSI and a fallback to the IPMSI. If the IPMSI next hops are not fully programmed at that moment, a temporary interruption of SPMSI flows occurs, while IPMSI flows continue forwarding normally. The traffic loss duration is short and typically limited to a few seconds to a few minutes during FPC upgrade cycles, not exceeding the expected ISSU convergence window. Both IPv4 and IPv6 multicast traffic are affected. PR1908581

Junos OS Release Notes for NFX Series

IN THIS SECTION

What's New

IN THIS SECTION

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

Additional Features

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

Support for Quantum Buffer in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptographic agility.

[See The Quantum Buffer and moduli.]

Support for Shor-resistant and other default key exchange algorithms in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSH supports the Shor's algorithm-resistant hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm. And, by default, SSH supports the `dh-group16-sha512` and `dh-group18-sha512` Diffie-Hellman (DH) group key exchange algorithms.

[See key-exchange.]

What's Changed

IN THIS SECTION

Learn about what's changed in this release for NFX Series.

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

EVPN

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit protocols evpn]` hierarchy level only on platforms that support a default switch instance. To enable

these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the [edit routing-instances name protocols evpn] hierarchy level.

[See Preference-Based DF Election.]

- **XML output tags changed for show l2-learning evpn mcnh-info command**—We've changed the XML output for this command. The XML <l2ald-mcnh-entry-list-headers> element level under the <l2ald-mcnh-entry> element has been removed. The child output tags that were under the <l2ald-mcnh-entry-list-headers> element are now direct child tags of the <l2ald-mcnh-entry-list> element.

General Routing

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the [edit services port-extender satellite <name> device-model] hierarchy, you can use the '?' to see a list of supported platforms. We have also added the skip-lo0-config statement under [edit service port-extender jnu] hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added cp-mtu under the [edit service port-extender satellite <name>] hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

  [See port-extender (Services).]

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for auto-configuration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address ff02::1. The RA source is the router's link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the show route extensive expanded-nh command.

  [See show route extensive <route> expanded-nh.]

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)               - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)        -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **Per-fan entries for fan trays in SNMP MIB tables (PTX10008)**—You can monitor individual fans in a fan tray through distinct entries in SNMP MIB tables. This enhancement addresses the earlier limitation that SNMP MIB does not list individual fans.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.

- A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Member Link Limit on Adaptive Aggregated Ethernet (PTX Series)**—A maximum of 64 member links are configurable on an aggregated Ethernet (AE) bundle with adaptive load balancing. Configuring more than 64 member links will result in a commit failure.

  [See Understanding Aggregated Ethernet Load Balancing.]

- Write cache disabled alarms no longer raised—The "Disk 2 Write Cache disabled" alarms were reported in the chassis alarms due to a missing SSD model in Junos OS supported device table. The support is now added, and alarms are no longer reported.

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading, copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

  [See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:

  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.

  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.

  - The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

[See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Network Management and Monitoring

- **Improved ZTP SNMP configuration handling (Junos)**—We have reordered the SNMP configuration commands to optimize the Zero Touch Provisioning (ZTP) process. The `set snmp engine-id local` statement now processes before `set snmp v3 usm local-engine user`, ensuring SNMPv3 USM authentication and privacy keys generation using the configured local engine ID instead of an auto-derived default. This prevents SNMPv3 query authentication failures from engine-ID resets.

  [See Configure Engine ID on SNMPv3 and Create SNMPv3 Users.]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **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.]

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **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.

User Interface and Configuration

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:

  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`

  - Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`

  - Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

  [See show system data-models.]

- **Introduction of remove operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents <userinput>statement not found</userinput> warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

- **Removal of weak algorithms in IPsec VPN with iked process (MX301, MX304, MX10004, MX10008, MX Series with MX-SPC3, NFX250, NFX350, SRX Series, and vSRX3.0)**—The following algorithms are no longer available as CLI options:

  - `des-cbc` and `3des-cbc` as encryption algorithms in IKE proposal and IPsec proposal at the `[edit security ike proposal encryption-algorithm]` and `[edit security ipsec proposal encryption-algorithm]` hierarchy levels respectively.

  - `md5` and `sha1` as authentication algorithms in IKE proposal at the `[edit security ike proposal authentication-algorithm]` hierarchy level.

  - `group1`, `group2`, and `group5` as DH groups in IKE proposal and Perfect Forward Secrecy (PFS) groups at the `[edit security ike proposal dh-group]` and `[edit security ipsec policy perfect-forward-secrecy keys]` hierarchy levels respectively.

  - `hmac-md5-96`, `hmac-sha1-96`, and `hmac-sha-256-96` as authentication algorithms in IPsec proposal at the `[edit security ipsec proposal authentication-algorithm]` hierarchy level.

  - `basic`, `compatible`, and `standard` as IKE proposal sets and IPsec proposal sets at the `[edit security ike policy proposal-set]` and `[edit security ipsec policy proposal-set]` hierarchy levels respectively.

  For MNHA HA link encryption and standalone VPNs, using `hmac-sha1-96` as the IPsec authentication algorithm at the `[edit security ipsec proposal proposal-name authentication-algorithm]` hierarchy level results in a commit error. Configure stronger algorithms to enhance IPsec VPN security.

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

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for NFX Series.

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

Routing Protocols

- larger primes of size greater than 8192 are generated when designer type moduli is confiugured with count 3 and 4

  PR1942246

Open Issues

IN THIS SECTION

Learn about open issues in this release for NFX Series.

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

General Routing

- NFX: "request vmhost reboot disk" doesn't work sometimes when one partition has pre-23.4R1 (Wind River 8 based) image version and the other has a 23.4R1 (Wind River LTS19 based) or later image version installed. PR1753117

- ON NFX150 HA When there is change in the vlan id from VRIOT to SRXFPE data path of interface on L2 interface we observe traffic drop for first few seconds. The behavior indicates that ARP and MAC learning occurred only after the 17th packet, which explains the initial packet loss following the VLAN change. Once ARP and MAC tables were populated, packets were forwarded normally without further loss. PR1867867

Resolved Issues

IN THIS SECTION

Learn about resolved issues in this release for NFX Series.

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

General Routing

- On all Junos NFX350 platform, the XE interface goes down after upgrading to Junos 24.2R2-S2. This behavior occurs when auto-negotiation is disabled on both sides, and the interfaces are manually configured for spped/duplex.

  PR1913204

Junos OS Release Notes for QFX Series

IN THIS SECTION

What's New

IN THIS SECTION

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

EVPN

Anycast multihoming for EVPN-VXLAN Type 2 symmetric IRB and BUM traffic (QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)—Use anycast multihoming with egress link protection (ELP) to achieve subsecond EVPN-VXLAN convergence for link and node failures with Type 2 symmetric IRB and Type 3 enhanced broadcast, unknown unicast, and multicast (BUM) traffic handling. This approach shifts load balancing to the underlay, reduces control-plane churn, and enables global repair. To configure anycast multihoming, use the following commands:

- `set protocols evpn anycast-multihoming [mlag-interop-mode | router-mac | inet/inet6 address | loopback-interface]`
- `set forwarding-options evpn-vxlan reroute-address`
- `set interfaces` *aeX* `esi anycast`
- (For BUM convergence) `set interfaces` *aeX* `esi df-election-granularity per-esi`

You can use Auto-ESI through Type 1 LACP. Constraints include no single-home support and no Type 5 anycast support.

Auto-detect UAP interfaces for Group-Based Policy (EX4000, EX4100, EX4400, EX4650, and QFX5120)—When configuring unified access policy, you no longer need to specify the UAP interface role. You can configure the switch to auto-detect whether the interface is connected to another switch or to a Mist AP.

Dynamic activation of VLANs with 802.1X in an EVPN-VXLAN network (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, 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 support dynamic activation of VLANs with 802.1X in an EVPN-VXLAN network. With this feature, VLANs are initially created locally on the Junos device in a deactivated (inactive) state. When a user client successful authenticates with 802.1X, the authentication server assigns the VLAN to the local port. Junos then activates the VLAN and installs the host routes in hardware.

To enable dynamic activation of VLANs with 802.1X, use `set vlans` *`vlan-name`* `activate-on-local-ports` or `set routing-instances` *`routing-instance-name`* `vlans` *`vlan-name`* `activate-on-local-port`.

[See Dynamic Activation of VLANs with 802.1X.]

EVPN maintenance mode scheduling (EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, QFX5120-32C, QFX5120-48T, and QFX5120-48Y)—You can schedule the EVPN maintenance mode CLI to minimize traffic loss and streamline upgrades across `erb-leaf`, `core-distribution`, `collapsed-core`, and `ip-clos` deployments. Use `set protocols evpn maintenance-mode (collapsed-core | core-distribution | erb-leaf | ip-clos) action-type (force | validate-and-set) schedule-start-at` *`time`* `[schedule-end-at` *`time`*`]` to trigger timed designated forwarder (DF) preference changes, network isolation (NISO), and Type 5 route withdrawal. Enable the NTP-based DF with `set protocols evpn df-election-ntp`. The `force` option overrides precheck results. Validate with `request evpn validate-maintenance-mode-pre-checks` *`deployment-mode`* and monitor with `show evpn maintenance-mode-status`. For Data Center Interconnect (DCI) gateways, delay interconnect Ethernet segment identifier (I-ESI) recovery using `set protocols evpn interconnect-esi-bringup-delay` *`seconds`*.

[See EVPN Maintenance Mode for Multihomed Leaf Isolation.]

EZ-LAG simplified CLI settings for native VLAN ID and access mode interfaces, and trunk mode interface configuration change (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, 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)—You can use the following new statements to define the peer provider edge (PE) device interfaces to the servers in an easy EVPN LAG (EZ-LAG) simplified configuration:

- Specify a user-defined native VLAN ID for trunk mode interfaces using the `native-vlan-id` *`vlan-id`* statement at the `[edit services evpn server` *`server-name`*`]` hierarchy level. You must include the existing `vlan-id-list [`*`vlan-id-list ...`*`]` statement to generate a trunk mode interface configuration.

- Generate PE-to-server interface configurations in access mode instead of trunk mode using the `access-vlan-id` *`access-vlan-id`* statement instead of the `vlan-id-list` statement. You can set `access-vlan-id`at the `[edit services evpn server` *`server-name`*`]` or `[edit services evpn server` *`server-name`* `mac-vrf-instance` *`instance-id`*`]` hierarchy levels.

Also, the EZ-LAG commit script now generates trunk interface configurations with the `flexible-vlan-tagging` option instead of the `vlan-tagging` option.

[See https://www.juniper.net/documentation/us/en/software/junos/evpn/topics/topic-map/easy-evpn-lag-config-script#concept_example-native-vlan-access-mode.html, server (Easy EVPN LAG Configuration), and mac-vrf-instance (Easy EVPN LAG Configuration).]

Global AS override for auto-derived route targets in EVPN (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, 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)—EVPN now supports global autonomous system (AS) override for auto-derived route targets. You can configure independent AS control for import and export route target generation by using the **import-as** and **export-as** statements. Alternatively, you can configure unified AS control for both import and export route target generation by using the **as-num** statement. This feature provides additional control over the AS value used for EVPN auto-derived route target generation in multi-AS deployments.

[See Global AS Override for Auto-Derived Route Targets.]

Junos Telemetry

Support for genstate YANG data models (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Junos telemetry supports the genstate YANG data model, a YANG-based model autogenerated from operational state data. Genstate models expose a subset of `show` command

outputs through gNMI subscriptions. This feature allows a gNMI telemetry collector to subscribe directly to specific resource paths in the genstate model to retrieve the required state data. The genstate models mirror selected hierarchies and fields from Junos `show` commands, allowing you to query information such as interface attributes, protocol status, system metrics, and more.

For example, you can retrieve the `evpn-instance` information by configuring the gNMI collector to subscribe to the following resource path:

```
/genstate/evpn-instance-information/evpn-instance
```

In this path, `evpn-instance` is a field displayed in the output of the `show evpn instance` command.

Genstate resource paths are prefixed with `/genstate/`. The genstate YANG files are published on Github. See Genstate YANG Data Models to view the supported show commands and the root paths.

[See Model-Driven Telemetry, Data Models, and Explore Sensor Paths.]

Decoupled rendering for telemetry data streaming (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—The Junos telemetry infrastructure now separates data rendering from core application processing, enhancing system performance and scalability. Dedicated rendering applications handle telemetry data rendering, enabling originating applications to focus on primary functions. This decoupling of functions improves efficiency, particularly during high-volume telemetry data streaming. Users may notice changes in telemetry producer names, as rendering is performed by the corresponding render applications. No configuration changes are needed, and existing subscriptions remain unaffected. This enhancement aligns the telemetry framework with the updated architecture, optimizing performance and scalability for telemetry data processing and streaming. See, Decoupled Rendering of Telemetry Data for a complete list of root sensor paths and respective new producer application names.

Network Management and Monitoring

Support for full-length SNMP engine ID (QFX5110, QFX5200, and QFX5210)—Use the `unique-id` *`hex-string`* option in the `engine-id` configuration statement to enable a full-length SNMP engine ID. The unique ID uses a hexadecimal string of the length between 5 and 32 bytes.

[See engine-id, Configure SNMP in Junos OS, and Configure Engine ID on SNMPv3.]

Routing Protocols

Support for interface-scoped TCP authentication for BGP unnumbered peers (QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, QFX10002, QFX10002-60C, QFX10008, and QFX10016)—We are extending support for interface-scoped authentication for BGP unnumbered peers in single-hop connectivity. The interface-scoped configurations enable TCP-MD5, TCP-Keychain, TCP AO, and Ron Bonica-draft authentication for BGP unnumbered peers using IPv6 link-local addresses. The authentication method is best suited for deployments with symmetric routing, where traffic for a TCP flow traverses the same logical interface in both directions. Authentication keys are associated directly with logical interfaces instead of IP addresses or prefixes as BGP unnumbered auto-discovered peering uses IPv6 link-local addresses. Authentication is applied based on the logical interface over which the BGP session is established and is independent of the dynamically learned peer address.

We also support nonstop active routing (NSR) support for interface-scoped authentication for BGP auto-discovered neighbors.

Note: When an authentication algorithm is specified, it must be configured with a key-chain and cannot be used with key.

[See https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/interace-scoped-authentication-bgp-unnumbered-peers.html.]

Software Installation and Upgrade

In-band ZTP in EVPN multihoming deployments (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, 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)—In-band zero-touch provisioning (ZTP) now supports EVPN multihoming deployments that use Layer 2 integrated routing and bridging (IRB) interfaces. In EVPN multihoming topologies, downstream devices can establish redundant connections to multiple upstream devices during Day 0 provisioning. Upstream devices participate in an EVPN-VXLAN fabric to provide Layer 3 onboarding connectivity through DHCP relay services.

[See In-band Zero Touch Provisioning.]

User Access and Authentication

Per-Server RADIUS Instances with Flexible Source and Routing Selection (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5140-24CD8O, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-64OD, QFX5241-64QD,

QFX5250-64OE, QFX5250-64OE-DOT2L, QFX5700, and QFX5700E)—This feature enables the configuration of multiple RADIUS authentication and accounting instances per server address, allowing each instance to use distinct source IP addresses and routing instances. It enhances flexibility in multi-VRF and segmented network environments by improving reachability and control over RADIUS traffic paths. The capability supports both UDP and TLS transports, ensuring secure and reliable communication while enabling better integration across virtual routing instances and complex network topologies.

[See Understanding RADIUS Profile-Based Server Instances for Multi-Instance Reachability.]

Support for OSPFv2 HMAC SHA-2 keychain authentication (EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, QFX5120-48Y, and QFX5120-48YM)—You can enable the OSPFv2 keychain module with HMAC-SHA2 authentication to authenticate packets reaching or originating from an OSPF interface. HMAC SHA2 algorithms include HMAC-SHA2-256, HMAC-SHA2-384, and HMAC-SHA2-512 as defined in *RFC 5709*. We also support these algorithms along with the HMAC-SHA2-224 algorithm. This feature ensures smooth transition from one key to another for OSPFv2 with enhanced security. We also support HMAC-SHA1 and HMAC-SHA2 authentication for virtual and sham links.

To enable OSPFv2 HMAC-SHA2 authentication, configure the `keychain` *`keychain-name`* statement at the `[edit protocols ospf area` *`area-id`* `interface` *`interface-name`* `authentication]` hierarchy level and `algorithm (hmac-sha2-224 | hmac-sha2-256 | hmac-sha2-384 | hmac-sha2-512)` option at the `[edit security authentication-key-chains key-chain` *`key-chain-name`*`]` hierarchy level. To enable keychain authentication support for OSPFv2 virtual links, configure the `keychain keychain-name` configuration statement at the `[edit protocols ospf area` *`area-id`* `virtual-link` *`neighbor-id`* `router-id transit-area` *`area-id`* `authentication]` hierarchy level. To enable keychains authentication support for OSPFv2 sham links, configure the **keychain keychain-name** configuration statement at the `[edit protocols ospf area` *`area-id`* `virtual-link` *`neighbor-id`* `router-id transit-area` *`area-id`* `authentication]` hierarchy level.

[See Understanding OSPFv2 Authentication.]

Additional Features

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

BGP Auto-discovery underlay in EVPN-VXLAN (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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48P, EX4400-48T, EX4650, QFX5120-48Y, and QFX5120-48Y-VC)—We support BGP auto-discovery underlay in EVPN-VXLAN networks.

[See BGP Auto-Discovered Neighbors and peer-auto-discovery.]

Flexible-match firewall filters with user-defined offsets and sFlow sampling action (QFX5200 and QFX5210)—Use firewall filters to match user-defined fields at Layers 2, 3, or 4 with custom byte offsets for precise traffic selection. Configure the `flexible-match-mask` (`match-start`, `byte-offset`, `bit-length`, `prefix`)

statement with a filter term such as `count` and an apply action such as `sampling`. Integrate with sFlow by configuring `sample-rate`, `collector`, and `interfaces` to export data for matched traffic. Supported families include `ethernet-switching`, `inet`, and `inet6`.

[See Filter-Based sFlow Packet Sampling Overview.]

Overlay and CE-IP ping and traceroute support with IPv6 underlay in EVPN-VXLAN (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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)—Overlay and CE-IP ping and traceroute support with IPv6 underlay in EVPN-VXLAN

[See ping overlay, ping, traceroute, and traceroute overlay.]

Support for hardware resource threshold monitoring for capacity planning (EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, QFX5120-48T-VC, and QFX5120-48Y-VC)—Hardware Resource Monitoring alerts you when a resource reaches an upper or lower threshold. This feature helps prevent traffic loss or device downtime through timely action.

[See Configure Hardware Resource Threshold Monitoring for Capacity Planning.]

Support for Quantum Buffer in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use

Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptographic agility.

[See The Quantum Buffer and moduli.]

Support for Shor-resistant and other default key exchange algorithms in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSH supports the Shor's algorithm-resistant hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm. And, by default, SSH supports the `dh-group16-sha512` and `dh-group18-sha512` Diffie-Hellman (DH) group key exchange algorithms.

[See key-exchange.]

Telemetry streaming for secure packet capture (EX4650, EX4650-48Y-VC, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, and QFX5120-48YM)—This feature enables you to capture packets from a device and securely stream them to an external telemetry data collector (for example, a cloud-based collector) for monitoring and analysis. Network professionals use real-time packet capture data to troubleshoot complex issues, including network and performance degradation, and poor end-user experience.

For secure packet capture, the maximum packet size is 512 bytes, including both the packet header and payload. To enable secure packet capture, include the */junos/system/linecard/packet-capture* resource path in a Junos RPC call.

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

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

This configuration sends packet capture sensor data to the collector. Remove the `packet-capture` configuration after the required data is collected. After capture is complete, ingress packets that match the filter conditions are trapped to the CPU and forwarded to the collector over a secure channel using the JTI-defined key-value format over a gRPC connection.

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 configure multiple interfaces for egress packet capture. When enabled, host-bound egress packets are captured and sent to the collector. Remove the configuration when packet capture is no longer required.

[See Junos YANG Data Model Explorer and Supported gRPC and gNMI Sensors.]

What's Changed

IN THIS SECTION

Learn about what's changed 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 OS Problem Report Search application.

EVPN

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit protocols evpn]` hierarchy level only on platforms that support a default switch instance. To enable these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the `[edit routing-instances name protocols evpn]` hierarchy level.

  [See Preference-Based DF Election.]

- **XML output tags changed for `show l2-learning evpn mcnh-info` command**—We've changed the XML output for this command. The XML `<l2ald-mcnh-entry-list-headers>` element level under the `<l2ald-mcnh-entry>` element has been removed. The child output tags that were under the `<l2ald-mcnh-entry-list-headers>` element are now direct child tags of the `<l2ald-mcnh-entry-list>` element.

General Routing

- **vxlan vni match condition enabled (EX4400-24MP, QFX5120-32C, QFX5120, QFX5120-32C)**—The following match conditions are now enabled on these platforms:

  - vxlan vni <vni-value>—Match the VNI.
  - vxlan vni-except <vni-value>—Do not match the VNI.

  [See Firewall Filter Match Conditions and Actions (PTX Series Routers).]

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the `[edit services port-extender satellite <name> device-model]` hierarchy, you can use the '?' to see a list of supported platforms. We have also added the `skip-lo0-config` statement under `[edit service port-extender jnu]` hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating

the lo0 configuration. We have added `cp-mtu` under the [`edit service port-extender satellite <name>`] hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

[See port-extender (Services).]

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for auto-configuration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address `ff02::1`. The RA source is the router's link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- **Sflow sampling with firewall filters**—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`with `setfirewall family inet filter filer-nameterm term-namethen 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 https://www.juniper.net/documentation/us/en/software/junos/network-mgmt/topics/concept/sflow-support-on-switches.html]

- **Enhanced loop detection for Layer 2 networks (EX4650, QFX5120-32C, QFX5120-48T, QFX5120-48Y, and QFX5120-48YM)** —We've added support for monitoring all VLANs on a logical interface. You can set `vlan-id all` option under the [`edit vlans` *`vlan-name`*] hierarchy. With this change, the devices can detect network loops across multiple VLANs and interfaces, improving network stability and performance.

  [See Enable Lightweight PE-CE Loop Detect on a Logical Interface and loop-detect.]

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the `show route extensive expanded-nh` command.

[See show route extensive <route> expanded-nh.]

- Updated documentation to show the new option: https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/statement/protocols-ptp.html https://www.juniper.net/documentation/us/en/software/junos/cli-reference/topics/ref/command/show-ptp-global-information-qfx-series.html

- **IFA metadata stack enhancement (QFX Series)**—Additional class-of-service fields are now encoded in the Inband Flow Analyzer (IFA) metadata. Along with the existing fields, three egress-queue related counters are included: Egress Queue Tx Bytes, Egress Queue Current Bytes (Current Buffer Occupancy), and Egress Queue Maximum Bytes (Peak Buffer Occupancy).

  [See Inband Flow Analyzer (IFA) 2.0 Real-Time Flow Monitoring.]

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)               - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)        -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- **Manual reboot required for specific forwarding and chassis options in clustered deployments (EX Series, QFX Series)**—For a Virtual Chassis or Virtual Chassis Fabric (VCF), the Packet Forwarding Engine in member switches does not automatically restart upon configuring and committing. When you commit any of the following statements, the commit succeeds but the new profile does not take effect until you manually reboot the system: `set forwarding-options vrf-fallback set forwarding-options mpls-tunnel-extend set forwarding-options vxlan-disable-copy-tos-encap set forwarding-options vxlan-disable-`

copy-tos-decap set chassis loopback-firewall-optimization set chassis per-logical-interface-firewall set chassis loopback-firewall-per-vrf After committing these configuration statements, you must schedule and perform a manual reboot for the changes to be applied to the forwarding plane on all members. Until you reboot, the system continues to use the existing forwarding profile.

- **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).]

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- Support added for new payload match conditions (EX4100 , EX4400 , 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 https://www.juniper.net/documentation/us/en/software/junos/routing-policy/topics/concept/firewall-filter-match-conditions-for-ipv4-traffic.html]

- [New output field added to the show system switchover command—The show system switchover command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **Per-fan entries for fan trays in SNMP MIB tables (PTX10008)**—You can monitor individual fans in a fan tray through distinct entries in SNMP MIB tables. This enhancement addresses the earlier limitation that SNMP MIB does not list individual fans.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

- The software has learned a new flow.
- A flow has aged-out.

Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

[See template (Inline Monitoring).]

- **Member Link Limit on Adaptive Aggregated Ethernet (PTX Series)**—A maximum of 64 member links are configurable on an aggregated Ethernet (AE) bundle with adaptive load balancing. Configuring more than 64 member links will result in a commit failure.

  [See Understanding Aggregated Ethernet Load Balancing.]

- Write cache disabled alarms no longer raised—The "Disk 2 Write Cache disabled" alarms were reported in the chassis alarms due to a missing SSD model in Junos OS supported device table. The support is now added, and alarms are no longer reported.

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading, copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

  [See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:
  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.
  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.

- The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

[See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Network Management and Monitoring

- **Improved ZTP SNMP configuration handling (Junos)**—We have reordered the SNMP configuration commands to optimize the Zero Touch Provisioning (ZTP) process. The `set snmp engine-id local` statement now processes before `set snmp v3 usm local-engine user`, ensuring SNMPv3 USM authentication and privacy keys generation using the configured local engine ID instead of an auto-derived default. This prevents SNMPv3 query authentication failures from engine-ID resets.

  [See Configure Engine ID on SNMPv3 and Create SNMPv3 Users.]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **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.]

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **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.

User Interface and Configuration

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:

  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`
  - Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`
  - Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

  [See show system data-models.]

- **Introduction of remove operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents <userinput>statement not found</userinput> warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

- **CLI options and RPC formats to retrieve `genstate` information and XPath expressions for operational commands**—A gNMI telemetry collector can subscribe to `genstate` resource paths to query for the corresponding state data on a device. We've added CLI pipe filter functions and RPC formats that enable you to discover available `genstate` resource paths and the associated `genstate` YANG module for a given show command. You can also retrieve the XPath expression for any node in command output. The options include:

  - `| display genstate subscription-paths`—Display all resource paths available for subscription in the `genstate` YANG module associated with the given command.
  - `| display genstate yang-module-name`—Display the name of the `genstate` YANG module that includes the data model and resource paths for the given command.
  - `| display xml genstate-path`—Display the command output in XML and include the `junos:genstate-path` attribute for every emitted tag that is available for subscription. The attribute value is the `genstate` subscription path. For RPCs, use the `format='xml-genstate-path'` attribute in the request tag.
  - `| display xml rpc response-paths`—List all possible response paths under the equivalent RPC's top-level response tag.
  - `| display xml xpath`—Display the command output in XML and include the `junos:data-xpath` attribute for each tag in the output. The attribute value is the XPath expression that selects that specific node. For RPCs, use the `format='xml-xpath'` attribute in the request tag.

  [See | (pipe) and Junos Genstate YANG Data Models.]

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

Known Limitations

IN THIS SECTION

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 OS Problem Report Search application.

General Routing

- For traffic dropped at egress due to Split Horizon in BCM during egress path processing, stats are shown on vtep as statistics are fetched at ingress pipeline in BCM! PR1656400

Routing Protocols

- larger primes of size greater than 8192 are generated when designer type moduli is confiugured with count 3 and 4

  PR1942246

Open Issues

IN THIS SECTION

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 OS Problem Report Search application.

General Routing

- QFX5210-64C: Packet loss is seen, when triggered Primary link flap PR1920865

- On QFX5K platforms, it is observed that due to a software issue, RE (Routing Engine) is consuming EAPOL (Extensible Authentication Protocol over LAN) packets on VXLAN (Virtual Extensible LAN) enabled interfaces even though layer2-control l2pt (Layer 2 Protocol Tunneling) is configured.

  PR1936048

- On QFX5210-64C platforms, the transceiver-related telemetry sensor paths do not function at the leaf level, as data retrieval from the leaf level does not returning any information. This results in a loss of telemetry data for leaf-level sensor paths.

  PR1926570

Interfaces and Chassis

- On some QFX platforms, configuring interface range using the wildcard et-0/0/* causes the system to incorrectly expand the range beyond the actual number of physical Ethernet ports. As a result, any subsequent configuration applied to the affected interface range fails to commit due to the presence of invalid interfaces.

  PR1853302

Resolved Issues

IN THIS SECTION

Learn about resolved 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 OS Problem Report Search application.

EVPN

- On Junos OS MX and QFX series platforms, in Virtual Router Redundancy Protocol (VRRP) with Ethernet Virtual Private Network - Virtual Extensible Local Area Network (EVPN-VXLAN) scenario,

when VRRP priority is changed, the VRRP virtual Media Access Control (MAC) address can remain pinned as a static entry on the remote device. As a result, MAC movement does not occur correctly, and VRRP packets are dropped on the leaf device. This impacts VRRP operation after priority changes.

PR1921796

- On QFX5210, QFX5220, QFX5210 - 64C platforms, inter-VLAN routed traffic fails, not because route targets aren't autoupdating, but due to missing or incorrect EVPNVxLAN (Ethernet Virtual Private Network - Virtual Extensible LAN) IRB(Integrated Routing and Bridging) symmetric routing configuration within the T5 VRF(Virtual routing and forwarding). PR1925825
- When Junos and Junos Evolved platform sends two proxied Neighbor Advertisement for default gateway, second one without Router flag and target Link-Layer address. The second Neighbor Advertisement can disrupt connectivity for end hosts. A proxied Neighbor Advertisement (NA) is an IPv6 mechanism where a network device responds to Neighbor Discovery requests on behalf of another device, instead of the actual endpoint replying itself. PR1926468

Forwarding and Sampling

- On specific VMHost platforms like MX204/MX304/MX10003/MX10008/MX10016/ PTX1000/ PTX10008/PTX10016/ EX2300/EX3400/EX4650/ QFX5100/QFX5110/QFX5200/ QFX10002/ QFX10008/QFX10016/ vSRX/SRX1500/SRX1600/SRX2300/SRX4100/SRX4200/SRX4300/ SRX4600 platforms. Whenever a client (e.g., lacp/ifinfo/snmp) requests statistics from PFE, the query is routed via the kernel and not through the use of the BULKGET protocol. PR1914977

General Routing

- On QFX5120-48Y, QFX5110-32q, QFX5110-48s platforms, running on 23.4R2-S4 release and multiple configurations are committed at the same time, the FPGA (Field-Programmable Gate Array) component responsible for managing QSFP (Quad Small Form Factor Pluggable) ports get into a hung state or reset state and the QSFP modules go undetected leading to loss of connectivity. PR1920870
- In an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) deployment, a remote MAC move event can cause the remote MAC IP entry to be removed after it is initially installed, resulting in missing Address Resolution Protocol (ARP) state and incomplete hardware programming on the remote Virtual Tunnel Endpoint (VTEP), which can lead to traffic loss for the affected endpoint. PR1936648
- l2ald Core @l2ald_gbptag_change_mac_in_db, @gbptagdb_update_internal_taginfo is seen in Suite teardown PR1925496
- On QFX5210 platform with Class of Service (CoS), the dcpfe process crashes when classifier is configured or removed on interface with active traffic. This can affect traffic forwarding till dcpfe process restarts post crash.

PR1921455

- On QFX5k platforms, if sflow is configured without multicast and BUM traffic is sampled, then the errors "fpc0 ifd null, port X and TD3:ifd null, port X" are observed. PR1932314
- While a filter with a routing-instance action was being reprogrammed as part of a CLI commit, a firewall VTY command to dump this filter was executed concurrently. This concurrent access resulted in a DC-PFE core, as filter-related data structures were modified during reprogramming. PR1924654
- On Junos OS QFX5210-64C platform, 100G optics inherently operate in CAUI-4 (Chip-to-Module 100 Gb/s Four-Lane Attachment Unit Interface) mode, which is automatically selected by the system and not user-configurable; mismatches due to software versions leading to link instability or flapping. PR1938876
- On all Junos EX and QFX platforms operating in an Ethernet Virtual Private Network (EVPN) environment, this issue occurs when the CLI command `licensing hourly update` is executed on a device that does not support the licensing feature.

PR1930036

- On all Junos OS QFX5K and EX4k platforms, configuring a native VLAN along with VLAN members on an underlay Integrated Routing and Bridging (IRB) Network-to-Network Interface (NNI) that carries VxLAN (Virtual Extensible LAN) traffic results in the VLAN tag not being stripped as expected. Instead of treating the native VLAN traffic as untagged, the interface adds the VLAN tag, leading to packet drops at the remote end. PR1939298
- On QFX5K and EX4K platforms which are running Junos release, when new VxLAN (Virtual Extensible LAN) vlans with IRBs (Integrated Routing & Bridging) are added, the L3 (Layer 3) interface values overwrite the MTU (Maximum Transmission Unit) entries previously programmed for non-VxLAN vlans that use the same hardware token internally. The VxLAN L3 interface is created with default MTU (1514) because L3 interface MTU value is still 1514 as it is not update by RE (Routing Engine). PR1905607
- The VSTP (VLAN Spanning Tree Protocol) bridge-priority configuration does not apply to VLAN (Virtual LAN) groups as expected. This issue results in VLAN groups not reflecting the intended bridge priority. This configuration issue can affect the network's root bridge election process, causing potential inconsistencies in network behavior and performance. PR1938991
- On Junos platforms having QFX and EX, VLAN(Virtual Local Area Network) traffic loss occur on interfaces configured with flexible-vlan-tagging when an L3 (SP style) sub-interface is deactivated on a port that also carries L2 (enterprise style) VLAN units. PR1938291
- Intermittent uplink ports flaps events observed and log messages show continuous syspld read failures during the issue. After a power cycle, the devices recovered. PR1925996

- On Junos OS QFX5120 and EX4650, when a L2 circuit is configured between two Customer Premises Equipment (CPE) , IPv6 Router Advertisement packets will not be sent to peer. This will cause devices to not exchange IPv6 messages so they cannot automatically be configured with an IPv6 address, therefore IPv6 connectivity won't be established. PR1922278

- On Junos QFX5110, QFX5120, EX4650, EX4400, EX4100 platforms, when operating with EVPN-VxLAN type5 overlay ECMP at high tunnel next hop scale, the Packet Forwarding Engine (PFE) crash when a memory corruption issue is observed due to a buffer overflow that leads to corruption of heap management structures. This corrupted metadata is detected by the memory manager when later heap allocation or free operations are invoked, resulting in a reported heap corruption condition.

  PR1933698

- On QFX5120, QFX5200, EX4400 platforms running Junos, when a firewall filter term includes the then sampling action and the user configures the sFlow sample rate only at the interface level using set protocols sflow interfaces <interface> sample-rate ingress <value>, the commit operation fails with the error: "sampling action required sflow sample-rate ingress and the sflow interface configured under protocols" PR1942162

- On Junos OS platforms, when the device is operating as a spine, with EVPN-VXLAN (Ethernet VPN Virtual Extensible LAN) in CRB (centrally routed bridging), enabling BFD (Bidirectional Forwarding Detection) as distributed mode on IRB (Integrated Routing and Bridging) interfaces, causes the BFD destination process to not to set the output interface correctly, preventing the BFD session from coming up. PR1948146

- On Junos OS QFX5k and EX4k platforms, executing the VTY command show firewall filter hw filter-name <non-existent-name> when a firewall filter is configured triggers a crash of the dcpfe process. This occurs when the command references a hardware filter name that does not exist. This results in a PFE (Packet Forwarding Engine) restart, leading to temporary forwarding plane traffic disruption while the process recovers.

  PR1944897

- On QFX5120-48YM platforms, proposal for upgrading EPDM version from 4.0.4 to latest 4.4.2

  PR1920896

Junos OS Release Notes for SRX Series

IN THIS SECTION

What's New

IN THIS SECTION

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

Application Security

HTTP/2 AppQoS based on DSCP with first-stream classification for consistent policy enforcement (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Enforce granular application QoS for HTTP/2 traffic using differentiated services code point (DSCP) based on first-stream classifications. The device applies the application classification of the first-stream to the entire HTTP/2 parent session and all subsequent streams. Packets in the session inherit a uniform AppQoS behavior derived from that initial classification. Using first-stream classification ensures that traffic is classified and prioritized by application type when multiple streams use different applications.

This approach integrates with traditional and unified policy frameworks, simplifying deployment without requiring additional configuration.

[See Application QoS.]

HTTP/2 inspection for IDP (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4120, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Intrusion detection and prevention (IDP) now inspects HTTP/2 traffic to enhance security on SRX Series Firewalls. IDP identifies HTTP/2 application traffic and applies relevant attack signatures during traffic inspection. This capability increases visibility into modern web application traffic and improves threat detection over HTTP/2 connections.

IDP inspection secures HTTP/2 traffic within connection sessions and stream sessions. It detects attacks and anomalies at the protocol level.

You can create custom signatures using HTTP/2 contexts such as frame length, frame type, flags, and stream ID.

[See Support for HTTP/2 Traffic Inspection.]

HTTP/2 inspection for SSL proxy (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSL proxy adds Application-Layer Protocol Negotiation (ALPN)-based HTTP/2 inspection with policy controls for TLS 1.2 and TLS 1.3. During the TLS handshake, the proxy negotiates the HTTP/2 identifier h2 with compliant peers and falls back to HTTP/1.1 when HTTP/2 is disabled or not allowed.

HTTP/2 is disabled by default. You can enable or disable it globally by using the *set services ssl proxy global-config http2 on|off* command. *D*isable HTTP/2 selectively per policy using the *set security policies from-zone <from-zone> to-zone <to-zone> policy <policy-name> then permit application-services ssl-proxy disable-http2* command.

The SSL proxy enforces protocols by stripping HTTP/3 ALPN, blocks prohibited TLS 1.2 cipher suites, rejects renegotiation, and terminates TLS 1.3 post handshake authentication.

[See SSL Proxy HTTP/2 Inspection with ALPN (TLS 1.2 and TLS 1.3).]

HTTP/2 protocol support (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—HTTP/2 protocol is a modern version of the HTTP protocol designed to improve application performance and efficiency by enabling multiple transactions over a single connection. Each connection logically divides into multiple streams, enabling your device to inspect, permit, or deny transactions independently while maintaining a single connection.

You can safely enable applications running over HTTP/2 without any additional configuration on your device.

[See HTTP/2 Protocol.]

HTTP/2 stream-based session logging enhancement (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, and SRX5800)—We've enhanced HTTP/2 session logging by introducing dedicated log types for stream sessions, eliminating the need to parse the session-type field and preventing incorrect session counts. The system generates **RT_FLOW_STREAM_*** logs for stream sessions while retaining **RT_FLOW_SESSION_*** logs for connection sessions.

This feature improves metric accuracy by using HTTP/2 frame size for byte counters and replaces packet counters with frame counters.

[See HTTP/2 Connection and Stream Session Logging.]

Support for cleartext HTTP/2 inspection (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Inspect unencrypted HTTP/2 traffic with cleartext HTTP/2 inspection. The firewall identifies HTTP/2 connections early, manages them as single TCP connections, and tracks each HTTP/2 stream independently. The system performs Layer 7 inspection and enforces policies on each stream, providing granular visibility and control over individual transactions within the same connection.

To configure cleartext inspection for HTTP/2, use the CLI commands: `set security application-services http2 plain-text` and `set services ssl proxy global config http2 on`. Verify your configuration with `show security application-services status` and `show security application-services monitoring`.

[See HTTP/2 Inspection: Support for Cleartext HTTP/2 Traffic.]

Support for HTTP/2 traffic in Juniper ATP Cloud (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)—Juniper ATP Cloud supports HTTP/2 traffic, giving you clear, stream-level visibility into HTTP/2 communications. It helps you detect threats accurately, inspect traffic deeply, and apply security policies consistently, without slowing down network performance. This support works seamlessly with Advanced Anti-Malware (AAMW), antivirus, Security Intelligence (SecIntel), and Security Metadata Streaming (SMS).

[See show services advanced-anti-malware statistics, show services anti-virus statistics, and show services security-metadata-streaming.]

Support for HTTP/2 in web filtering and content filtering (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—The HTTP/2 protocol supports Content Security features such as web filtering and content filtering. These features bypass the HTTP/2 parent session and process child sessions directly. The system applies actions based on the existing configuration for specific sessions, ensuring accurate filtering and enhanced security.

[See Content Security Overview and HTTP/2 Protocol.]

Content Security

Deprecation of performance-mode and disable-dynapp-profile-selection statements (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We've deprecated the `performance-mode` and `disable-dynapp-profile-selection` statements from the [`edit security utm default-configuration web-filtering`] hierarchy level.

[See web-filtering.]

Support for dynamic file type signatures in content filtering (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use dynamic file type signatures to efficiently manage file type signature data. These signatures are stored in a structured JSON format on a cloud-based server, allowing you to download, install, and update them without upgrading the existing Junos OS on your device. Before this release, you had to wait for the next Junos OS upgrade to update the file type signature data. Use CLI commands to easily manage tasks such as downloading, installing, and rolling back file type signatures.

[See request security utm content-filtering file-signatures download, request security utm content-filtering file-signatures download status, request security utm content-filtering file-signatures install, request security utm content-filtering file-signatures install status, request security utm content-filtering file-signatures local-update, request security utm content-filtering file-signatures rollback, request security utm content-filtering file-signatures rollback status, request security utm content-filtering file-signatures uninstall, request security utm content-filtering file-signatures uninstall status, and show security utm content-filtering file-signatures status.]

Flow-Based and Packet-Based Forwarding

Optimization of TCP Proxy (cSRX, SRX1600, SRX2300, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We can use SRX Series Firewall as a TCP Proxy. This configuration ensures consistent proxying and prevents implicit services from disrupting sessions. It enhances session reliability and security. To configure security policy with default (global TCP Profile), use the `set security policies from-zone trust to-zone untrust policy p1 then permit tcp-proxy` command.

To configure the TCP proxy profile, use the `set security flow tcp-proxy tcp-profile` command. This command helps in configuring the required parameters.

To verify the configuration, use the`show security flow tcp-proxy profiles` command.

[See TCP Sessions.]

High Availability (HA) and Resiliency

Domain-local IPsec VPN support in four-node MNHA deployments (SRX4700)—Four-node Multinode High Availability now supports multi SRG1 (SRG1+) within domains, enabling IPsec VPN in the deployment. In this setup, IPsec VPN tunnels are scoped per domain and do not span the entire deployment. Each domain operates as an independent two-node MNHA IPsec pair, and IKE/IPsec security associations synchronize only within that domain. This design eliminates cross-domain state sharing by scoping IPsec SAs per domain, ensuring predictable per-domain failover.

[See Four-Node and Three-Node Multinode High Availability.]

Dual-stack virtual IP addresses support for default gateway and hybrid deployments in Multinode High Availability (SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We support dual-stack (IPv4 and IPv6) addresses for virtual IP for Multinode High Availability (MNHA) operating in default gateway mode or hybrid mode. The `virtual-ip` configuration under the `[edit chassis high-availability services-redundancy-group]` hierarchy is extended to accept a list of IP addresses, allowing both IPv4 and IPv6 virtual IP address configuration on the same interface.

MNHA uses both address families to detect and resolve split brain conditions. If either address family is reachable, the probe succeeds and the services redundancy group state machine responds accordingly. This support improves MNHA adoption in dual stack environments.

[See Example: Configure Multinode High Availability in a Default Gateway Deployment and IP Addressing and Network Preparation for Multinode High Availability.]

Interfaces

PPPoE support (SRX1600)—Use Point-to-Point Protocol over Ethernet (PPPoE) to connect multiple Ethernet LAN hosts to a remote site through a single customer premises equipment (CPE) device. PPPoE aggregates access so hosts share a common DSL, cable modem, or wireless link to the Internet.

[See Configure Point-to-Point Protocol over Ethernet (PPPoE).]

Junos Telemetry

Support for genstate YANG data models (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Junos telemetry supports the genstate YANG data model, a YANG-based model autogenerated from operational state data. Genstate models expose a subset of `show` command outputs through gNMI subscriptions. This feature allows a gNMI telemetry collector to subscribe directly to specific resource paths in the genstate model to retrieve the required state data. The genstate models mirror selected hierarchies and fields from Junos `show` commands, allowing you to query information such as interface attributes, protocol status, system metrics, and more.

For example, you can retrieve the `evpn-instance` information by configuring the gNMI collector to subscribe to the following resource path:

```
/genstate/evpn-instance-information/evpn-instance
```

In this path, `evpn-instance` is a field displayed in the output of the `show evpn instance` command.

Genstate resource paths are prefixed with `/genstate/`. The genstate YANG files are published on Github. See Genstate YANG Data Models to view the supported show commands and the root paths.

[See Model-Driven Telemetry, Data Models, and Explore Sensor Paths.]

Decoupled rendering for telemetry data streaming (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX5500, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—The Junos telemetry infrastructure now separates data rendering from core application processing, enhancing system performance and scalability. Dedicated rendering applications handle telemetry data rendering, enabling originating applications to focus on primary functions. This decoupling of functions improves efficiency, particularly during high-volume telemetry data streaming.

Users may notice changes in telemetry producer names, as rendering is performed by the corresponding render applications. No configuration changes are needed, and existing subscriptions remain unaffected. This enhancement aligns the telemetry framework with the updated architecture, optimizing performance and scalability for telemetry data processing and streaming. See, Decoupled Rendering of Telemetry Data for a complete list of root sensor paths and respective new producer application names.

Junos XML API and Scripting

Event policy archive sites support specifying a routing instance (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When you configure an archive site for an event policy destination, you can specify the routing instance that the device uses to connect to the site. To configure the routing instance, include the `routing-instance` option at the `[edit event-options destinations` *`destination-name`* `archive-sites` *`url`*`]` hierarchy level.

[See Event Policy File Archiving and archive-sites (Event Policy).]

MACsec

Support for MACsec features over WAN (SRX4300 and SRX4700)—Configure Media Access Control Security (MACsec) on logical interfaces to extend hop-to-hop MACsec security to point-to-point connections. Combined, the MACsec features supported on logical interfaces allow you to establish more secure VLAN-level MACsec connections in enterprise WAN and service provider networks. These devices support the following MACsec features on logical interfaces (IFLs) in routing mode for Layer 3 traffic:

- Custom EAPoL destination MAC address for unicast, PAE, provider bridge, and LLDP multicast
- Sending single-tagged VLAN IDs in clear text to support VLAN-level MACsec
- GCM-AES-128, GCM-AES-256, GCM-AES-XPN-128, GCM-AES-XPN-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

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.

View MACsec statistics using the `show security macsec statistics interface` *`ifl-name`* and `show interfaces statistics` *`ifl-name`* `(detail | extensive)` commands.

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

Network Address Translation (NAT)

Deterministic NAT with interim logging (cSRX, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—You can now configure syslog intervals for deterministic NAPT. The configurable syslog interval reduces logging requirements by ensuring consistent IP address and port mappings.

Use the `set security nat source pool <pool-name> port deterministic det-nat-configuration-log-interval` command to define intervals from 1800 through 86400 seconds.

[See Deterministic NAT and det-nat-configuration-log-interval.]

Dynamic scaling for persistent NAT capacity (cSRX, SRX1600, SRX2300, SRX4120, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—You can enhance persistent NAT capacity by dynamically scaling beyond the fixed, code-defined limit. Dynamic scaling in persistent NAT enables growth by up to a configurable number of additional entries, depending on the configuration and available runtime resources, in addition to the minimum baseline static bindings allowed on the device.

You must configure the `set security nat source persistent-nat increased-scale-limit <value>` command to increase the persistent NAT capacity.

[See Persistent NAT and NAT64, show security nat source persistent-nat-table, and increased-scale-limit.]

Network Management and Monitoring

NDP Proxy User Route (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, and SRX5800)—Use NDP proxy user route to proxy Neighbor Discovery Protocol (NDP) requests for user routes. SRX Series Firewalls maintain IPv6 user-route reachability by responding to NDP requests on behalf of reachable user routes. This capability benefits IPv6 subscriber or user-route deployments by eliminating

the need for endpoints to respond directly to NDP requests. It simplifies neighbor discovery behavior and improves IPv6 reachability handling in supported SRX Series Firewalls deployments.

[See NDP Proxy and DAD Proxy.]

On-box security log database capacity enhancement (SRX4600 and SRX4700)—We've increased the on-box security log database capacity to one billion entries on the SRX4600 (previously 512 million) and two billion entries on the SRX4700 across root and logical systems. Default database entries are 30 million/8 million (session/non-session tables) for the SRX4600 and 60 million/24 million for the SRX4700. Maximum entries are 110 million/65 million and 220 million/130 million, respectively. Configure the maximum capacity with the `set security log report database-capacity max-db-size` command and adjust per-table allocation with `set security log report log-share-per-table`.

[See Syntax (SRX4600 and SRX4700) and Platform-Specific System Logging Behavior.]

Services Applications

Support for GRE key insertion and tunnel multiplexing (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1600, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, and SRX5800)—GRE key support on SRX firewalls introduces optional GRE header key validation to strengthen tunnel correctness and security while preserving existing virtual routing and forwarding (VRF)-based demultiplexing behavior. The device inserts the key in outbound GRE packets and validates or parses keys from incoming GRE packets.

[See GRE Tunnel Key Support.]

User Access and Authentication

Sequential DNS fallback for outbound SSH (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When a configured hostname resolves to multiple IP addresses via DNS, Junos OS automatically tries each address in sequence, regardless of IP family, until a connection succeeds or all addresses have been attempted. The `outbound-ssh` retry mechanism activates after all resolved addresses have been tried.

[See outbound-ssh.]

VPNs

Configurable PRF in IKEv2 with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—You can configure Internet Key Exchange version 2 (IKEv2) proposals with a specific pseudorandom function (PRF) algorithm independent of the authentication algorithm. Use the `prf-algorithm` option to configure the following algorithms at the [`edit security ike proposal` *`proposal-name`*] hierarchy level in the IPsec VPN service with the iked process:

- `prf-hmac-sha-256`
- `prf-hmac-sha-384`
- `prf-hmac-sha-512`
- `prf-hmac-sha1`

You cannot combine a PRF with Authenticated Encryption with Associated Data (AEAD) cipher suites such as AES-GCM or ChaCha20-Poly1305 algorithm.

[See prf-algorithm and proposal (Security IKE).]

Deprecation of certificates as authentication method in IKE (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We've deprecated support for the `certificates` option at the [`edit security ike proposal` *`proposal-name`* `authentication-method`] hierarchy level. However, you can continue to use the other authentication methods at this hierarchy level.

[See proposal (Security IKE).]

Multiple local certificates in IKE with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—You can specify multiple local certificates in an IKE policy on a device that runs the iked process for the IPsec VPN service. During IKE negotiation, the device selects the local certificate whose certificate chain matches the peer's certificate request. The device evaluates the certificates in their listed order. Use `set security ike policy` *`policy-name`* `certificate local-certificate [`*`cert1 cert2`*`]` to configure multiple local certificates.

[See certificate and IKE Policy with Local Certificate.]

RSA-PSS signature algorithm in IKEv2 with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Specify the RSA signature scheme for IKEv2 authentication with an RSA local certificate as per RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2*. Configure the following algorithms using the `digital-signature-scheme` option at the [`edit security ike proposal` *`proposal-name`*] hierarchy level:

- `rsassa-pkcs1-v1-5` for RSA signature with the PKCS1 v1.5 scheme
- `rsassa-pss` for RSA signature with Probabilistic Signature Scheme (PSS) as the default signature scheme for RSA certificate

Do not configure `digital-signature-scheme` without setting `digital-signature` as the authentication method.

[See digital-signature-scheme, proposal (Security IKE), and show security ike security-associations.]

Support for trusted CA group for certificate validation in IKE with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Junos OS uses the trusted certification authority (CA) group in an IKE policy to validate peer certificates during IKE authentication. You can use the feature on a device that runs the iked process for the IPsec VPN service. Configure the trusted CA group to specify which CAs the device trusts for incoming peer certificate validation. Use `set security ike policy` *`policy-name`* `certificate trusted-ca trusted-ca-group` *`trusted-ca-group-name`* to reference the trusted CA group in an IKE policy for certificate validation.

[See certificate and IKE Policy with Trusted CA Group.]

Additional Features

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

Data path trace debugging (SRX4700)—We've enhanced the data path trace debugging feature to improve flow troubleshooting efficiency by providing:

- Packet-level visibility—Tracks packets across the data path, including ingress, policy lookup, NAT, and egress processing stages.
- Missing log detection—Highlights absent trace logs, helping identify potential packet drops or misconfigurations.
- Controlled trace capture—Limits output to a defined set of initial trace entries, reducing log noise while preserving essential context.

[See Configure Data Path Debugging and Trace Options.]

Encryption with TPM 2.0 (SRX1600, SRX2300, SRX4120, SRX4300, and SRX4700)—Use Trusted Platform Module (TPM) 2.0 to secure sensitive data including the master password on SRX Series devices.

[See Master Password for Configuration Encryption.]

Support for Quantum Buffer in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP,

EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptographic agility.

[See The Quantum Buffer and moduli.]

Support for Shor-resistant and other default key exchange algorithms in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSH supports the Shor's algorithm-resistant hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm. And, by default, SSH supports the `dh-group16-sha512` and `dh-group18-sha512` Diffie-Hellman (DH) group key exchange algorithms.

[See key-exchange.]

Supported transceivers, optical interfaces, and DAC cables (MX240, MX301, MX480, MX960, SRX4600, SRX5400, SRX5600, and SRX5800)—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.

What's Changed

IN THIS SECTION

Learn about what's changed in this release for SRX Series.

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

Chassis Clustering

- IPv6 support for Multinode High Availability Configuration —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]

- MNHA Upgrade Requirement for IPv6-Based ICL Encryption [SRX4600, SRX4200, SRX4300, 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 {https://supportportal.juniper.net/s/article/SRX-Isolated-Node-Upgrade-Procedure-for-Multinode-High-Availability} Isolated Node Upgrade Procedure for Multinode High Availability.] [For Software Upgrade in Multinode High Availability, see {https://www.juniper.net/documentation/us/en/software/junos/high-availability/topics/topic-map/mnha-upgrade-process.html} Software Upgrade in Multinode High Availability.]

EVPN

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit protocols evpn]` hierarchy level only on platforms that support a default switch instance. To enable these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the `[edit routing-instances name protocols evpn]` hierarchy level.

  [See Preference-Based DF Election.]

- **XML output tags changed for `show l2-learning evpn mcnh-info` command**—We've changed the XML output for this command. The XML `<l2ald-mcnh-entry-list-headers>` element level under the `<l2ald-mcnh-entry>` element has been removed. The child output tags that were under the `<l2ald-mcnh-entry-list-headers>` element are now direct child tags of the `<l2ald-mcnh-entry-list>` element.

General Routing

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the `[edit services port-extender satellite <name> device-model]` hierarchy, you can use the '?' to see a list of supported platforms. We have also added the `skip-lo0-config` statement under `[edit service port-extender jnu]` hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added `cp-mtu` under the `[edit service port-extender satellite <name>]` hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

[See port-extender (Services).]

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for auto-configuration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address `ff02::1`. The RA source is the router's link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the `show route extensive expanded-nh` command.

  [See show route extensive <route> expanded-nh.]

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)               - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)        -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- **SSL proxy protocol-version deprecation and new all-secure option (SRX Series, vSRX)**—TLS 1.0 and TLS 1.1 protocol versions are deprecated for SSL proxy-rather than immediately removed-to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration. As a part of this change, the corresponding protocol-version configuration options `tls1`, `tls11`, and `tls12-and-lower` under [`edit services ssl termination profile profile-name`], [`edit services ssl initiation profile profile-name`], and [`edit services ssl proxy profile profile-name`] are no longer available as standard options. A new `all-secure protocol-version` option is introduced to allow you to enforce only secure TLS versions (TLS 1.2 and TLS 1.3) for SSL proxy traffic.

  [See protocol-version.]

- **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.**

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- **Express Path support for ESP transit and ICL traffic (SRX4600, SRX4700, and SRX5800)**—Enhanced support for Express Path for Encapsulating Security Payload (ESP) transit and interchassis Link (ICL) traffic. With this enhancement, the Services Processing Unit (SPU) processes the first packet of an ESP session and offloads subsequent packets to the IOC or Network Processing Unit (NPU), improving forwarding performance and scalability. As a result, screen features such as `alarm-without-drop` for IPv6 extension header attacks are triggered only for the first packet in an ESP session, while subsequent packets bypass SPU processing. This behavior is expected when Express Path is enabled.

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **Per-fan entries for fan trays in SNMP MIB tables (PTX10008)**—You can monitor individual fans in a fan tray through distinct entries in SNMP MIB tables. This enhancement addresses the earlier limitation that SNMP MIB does not list individual fans.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.

  - A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Member Link Limit on Adaptive Aggregated Ethernet (PTX Series)**—A maximum of 64 member links are configurable on an aggregated Ethernet (AE) bundle with adaptive load balancing. Configuring more than 64 member links will result in a commit failure.

  [See Understanding Aggregated Ethernet Load Balancing.]

- Write cache disabled alarms no longer raised—The "Disk 2 Write Cache disabled" alarms were reported in the chassis alarms due to a missing SSD model in Junos OS supported device table. The support is now added, and alarms are no longer reported.

Intrusion Detection and Prevention (IDP)

- **Support for best effort snort rule conversion (SRX Series)**—A new `best-effort-conversion` option has been added to the `request security idp jist-conversion` command. This option allows the conversion process to continue by skipping unsupported Snort modifiers instead of stopping the conversion.

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading, copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

  [See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:

  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.

  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.

  - The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

  These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

  [See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Network Management and Monitoring

- **Improved Syslog User Identification (SRX4300)**—Syslogs generated from shell commands now logs the remote username instead of the local template username when you use the RADIUS or TACACS + authentication.

- **Improved ZTP SNMP configuration handling (Junos)**—We have reordered the SNMP configuration commands to optimize the Zero Touch Provisioning (ZTP) process. The `set snmp engine-id local` statement now processes before `set snmp v3 usm local-engine user`, ensuring SNMPv3 USM authentication and privacy keys generation using the configured local engine ID instead of an auto-derived default. This prevents SNMPv3 query authentication failures from engine-ID resets.

  [See Configure Engine ID on SNMPv3 and Create SNMPv3 Users.]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **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.]

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

- **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.

Unified Threat Management (UTM)

- **Support for IPv6 in Enhanced Web filtering (EWF) (SRX Series)**—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.

- 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.]

- **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.]

- **Enhanced anti-virus and web-filtering operational command outputs (SRX Series Firewall and vSRX)**—We've enhanced the `show security utm web-filtering category status` and `show security utm anti-virus status` outputs. The enhanced web-filtering output displays error message for the failed category-package file download. When you enable Avira engine, the enhanced anti-virus output displays error messages for the failed Avira pattern-updater download.

  Proxy profile settings update (host, port, username, password) causes download issues or connectivity problems, which are accurately reflected in the anti-virus and web-filtering operational command outputs.

  [See show security utm web-filtering category status and show security utm anti-virus status.]

User Interface and Configuration

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:

  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`
  - Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`
  - Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

  [See show system data-models.]

- **Introduction of remove operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents <userinput>statement not found</userinput> warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

- **CLI options and RPC formats to retrieve `genstate` information and XPath expressions for operational commands**—A gNMI telemetry collector can subscribe to `genstate` resource paths to query for the corresponding state data on a device. We've added CLI pipe filter functions and RPC formats that enable you to discover available `genstate` resource paths and the associated `genstate` YANG module for a given show command. You can also retrieve the XPath expression for any node in command output. The options include:

  - `| display genstate subscription-paths`—Display all resource paths available for subscription in the `genstate` YANG module associated with the given command.
  - `| display genstate yang-module-name`—Display the name of the `genstate` YANG module that includes the data model and resource paths for the given command.
  - `| display xml genstate-path`—Display the command output in XML and include the `junos:genstate-path` attribute for every emitted tag that is available for subscription. The attribute value is the `genstate` subscription path. For RPCs, use the `format='xml-genstate-path'` attribute in the request tag.
  - `| display xml rpc response-paths`—List all possible response paths under the equivalent RPC's top-level response tag.

- `| display xml xpath`—Display the command output in XML and include the `junos:data-xpath` attribute for each tag in the output. The attribute value is the XPath expression that selects that specific node. For RPCs, use the `format='xml-xpath'` attribute in the request tag.

[See | (pipe) and Junos Genstate YANG Data Models.]

VPNs

- **Configuration validation for HA link encryption (SRX4600, SRX4300, and SRX4700)**—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 where tunnel-mtu is less than 2000, you'll receive a commit check error:

  ```
  user@host# set security ipsec vpn L3HA_IPSEC_VPN tunnel-mtu <bytes>
  ```

  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 Overview.]

- SRX4700 does not support policy-based IPSec VPNs with IKEv2. Support for policy-based IPsec VPN is not available when using junos-ike package with your firewall running iked process for IPsec VPN service. With junos-ike package, remove any policy-based IPsec VPN configurations as they are ineffective.

- **Secure tunnel (st0.x) interface state behavior (SRX Series and vSRX 3.0)**—Previously, when `establish-tunnel immediately` was configured on a point-to-point st0.x interface, the interface followed the VPN tunnel state. For a dynamic peer VPN tunnel, the device converted `establish-tunnel immediately` to `establish-tunnel responder-only`, making the interface to stay up regardless of the tunnel state. With this release, point-to-point st0 interfaces consistently follow the VPN tunnel state (up or down) with `establish tunnel immediately`, `establish tunnel responder only`, and `establish tunnel responder only no rekey`, including for dynamic peers that use `establish-tunnel responder-only` or `establish-tunnel responder-only-no-rekey`. If multiple IPsec security associations (SAs) can establish from a single VPN tunnel configuration, the interface goes up when the first SA is established and goes down when the last SA is removed.

  [See IPsec VPN User Guide.]

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

[See show route extensive.]

- **Removal of weak algorithms in IPsec VPN with iked process (MX301, MX304, MX10004, MX10008, MX Series with MX-SPC3, NFX250, NFX350, SRX Series, and vSRX3.0)**—The following algorithms are no longer available as CLI options:

  - `des-cbc` and `3des-cbc` as encryption algorithms in IKE proposal and IPsec proposal at the `[edit security ike proposal encryption-algorithm]` and `[edit security ipsec proposal encryption-algorithm]` hierarchy levels respectively.

  - `md5` and `sha1` as authentication algorithms in IKE proposal at the `[edit security ike proposal authentication-algorithm]` hierarchy level.

  - `group1`, `group2`, and `group5` as DH groups in IKE proposal and Perfect Forward Secrecy (PFS) groups at the `[edit security ike proposal dh-group]` and `[edit security ipsec policy perfect-forward-secrecy keys]` hierarchy levels respectively.

  - `hmac-md5-96`, `hmac-sha1-96`, and `hmac-sha-256-96` as authentication algorithms in IPsec proposal at the `[edit security ipsec proposal authentication-algorithm]` hierarchy level.

  - `basic`, `compatible`, and `standard` as IKE proposal sets and IPsec proposal sets at the `[edit security ike policy proposal-set]` and `[edit security ipsec policy proposal-set]` hierarchy levels respectively.

  For MNHA HA link encryption and standalone VPNs, using `hmac-sha1-96` as the IPsec authentication algorithm at the `[edit security ipsec proposal proposal-name authentication-algorithm]` hierarchy level results in a commit error. Configure stronger algorithms to enhance IPsec VPN security.

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

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for SRX Series.

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

J-Web

- J-Web becomes unreachable when a local self-signed certificate with weak digital signature is assigned to the HTTPS PKI configuration. PR1955001

Routing Protocols

- larger primes of size greater than 8192 are generated when designer type moduli is confiugured with count 3 and 4

  PR1942246

Open Issues

IN THIS SECTION

Learn about open issues in this release for SRX Series.

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

General Routing

- On Junos SRX4700, SRX5400, SRX5600 and SRX5800 platforms with services offloading (SOF) enabled along with AppTrack and secure-wire configuration, IP TTL value is decremented for the packet passing through the secure-wire.

  PR1952647

Intrusion Detection and Prevention (IDP)

- IDP: On-box packet capture fails when packet-log host or port is not reachable with SSL Initiation Profile For On-box packet capture, ssl initiation profile configuration is NOT required. This issue will not be seen if we do not configure ssl initiation profile. PR1921354

Routing Policy and Firewall Filters

- On SRX5K platforms, during policy configuration processing using the policy file serialization feature, the Network Security Daemon (NSD) main thread can stop yielding for more than one second. This results in warning messages in system logs and can impact protocol/FPC stability if the delay becomes excessive.

  PR1932245

Resolved Issues

IN THIS SECTION

Learn about resolved issues in this release for SRX Series.

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

Application Layer Gateways (ALGs)

- On MX platforms with SPC3 line-cards/ SRX platforms, when there are bandwidth attributes carried in Session Description Protocol (SDP) messages inside Session Initiation Protocol(SIP) packets, SIP application traffic and SIP calls get impacted.

  PR1930403

General Routing

- On SRX 3XX platforms, due to incorrect mbuf (Memory Buffer) allocation, the ISIS (Intermediate System-to-Intermediate System) packets with MPLS (Multiprotocol Label Switching) tag transiting over a L2VPN (Layer 2 Virtual Private Network) are dropped if the packet size exceeds 1500 bytes resulting in the ISIS adjacency remaining stuck in the initializing state, leading to routing issues. This issue is intermittent and observed only when the SRX 3XX device is serving as the PE in MPLS L2VPN network. PR1887306

- On SRX3xx series platforms, when the LTE (Long Term Evolution) configuration is deleted and added back, device will reboot. PR1913687

- MAC limit Count show 65k in srx1600, but it stops learning at 28k. PR1916697

- When AE interface (member links) are configured with MACsec, the MKA session establishes successfully when the default standard MAC address is used. However, when a custom unicast MAC address is configured, the MKA session does not come up. This is not a common use case, but the workaround is to allow the MKA session to operate with the default MAC address instead of a custom unicast MAC address. PR1909930

- After reboot on branch SRX300 Series platforms, interfaces using Small Formfactor Pluggable (SFP) SX/LX fiber transceivers and configured with 'family ethernet-switching' and 'no auto negotiation' might remain in a down state due to incorrect physical medium detection by the Ethernet physical layer (PHY). The issue is timing-dependent, it might happen with single or after several reboots. PR1927646

- On all Junos platforms with PPPoE (Point-to-Point Protocol over Ethernet) interface, IP address assignment fails when PPPoE underlying interface from a higher-numbered port is changed to a lower-numbered port (for example, from ge-0/0/2 to ge-0/0/1) and this will lead to traffic drop. PR1918920

- On SRX300-series devices, when ethernet-switching is used with spanning-tree protocol enabled, in some cases ARP reply packets may be sent out from a spanning-tree blocked port. PR1913911

- On the SRX380 platforms, in both cluster and standalone modes, the XE interfaces with 1G SFP fiber transceivers flap once after every configuration commit, causing the device to temporarily lose upstream reachability. PR1934517

- On SRX1500 platform, when an SFP-T transceiver port (ge-0/0/12 to ge-0/0/15) is administratively disabled, the local interface shows down, but the link on the peer device remains up. The link status of a disabled SFP-T port on SRX1500 is down, but peer remote port remains up.

  PR1932505

J-Web

- When a MAC-RADIUS authentication succeeds and the RADIUS server returns CWA redirect attributes (URL-Redirect), the EX switch should intercept client HTTP traffic and return an HTTP 302 redirect to the Captive Portal. The switch instead responded with HTTP 405 (Method Not Allowed), causing the client to bypass the CWA redirect workflow. PR1931780

Routing Policy and Firewall Filters

- On SRX devices operating in a High Availability cluster with logical-systems enabled, the backup node may display fewer flow sessions than the primary node due to incomplete flow session synchronization. Sessions that do not synchronize to the backup node are lost during a failover, resulting in service interruption for the affected flows.

  PR1939433

Unified Threat Management (UTM)

- On Junos OS SRX platforms, configuring cache preload in web filtering causes high memory utilization in the UTM (Unified Threat Management) pool, which results in traffic loss. PR1927484

VPNs

- On SRX series platforms using the IPsec Key Management Daemon (KMD), RG0 failover or failback while Auto Discovery VPN (ADVPN) shortcuts are active may cause the KMD process to crash, temporarily disrupting ADVPN sessions.

  PR1922670

- On all SRX platforms, when interface st0 inet6 is initially configured as point-to-point, and later if st0 is modified to multi-point along with the configuration of IKE (Internet Key Exchange)/IPsec (Internet Protocol Security), which uses this st0, the ikemd enters a timing issue where it can't re-read the LLA (Link-Local Address) information of this st0 while parsing ike/ipsec configuration. Thus, resulting in a situation where the NHTB (Next-Hop Tunnel Binding) table is not complete, and the OSPF(Open Shortest Path First) IPv6 neighbor discovery fails. PR1918367

Junos OS Release Notes for vRR Series

IN THIS SECTION

What's New

IN THIS SECTION

Learn about new features introduced in this release for vRR Series.

Junos Telemetry

Enhancement of BGP XPath policy (ACX5448, MX480, MX960, MX2020, MX10004, and vRR)—Add OpenConfig BGP XPath enhancements for default policy and policy definitions.

[See Junos YANG Data Model Explorer.]

What's Changed

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

Known Limitations

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

Open Issues

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

Resolved Issues

There are no resolved issues in this release for vRR Series.

Junos OS Release Notes for vSRX

IN THIS SECTION

What's New

IN THIS SECTION

Learn about new features introduced in this release for vSRX.

Application Security

HTTP/2 AppQoS based on DSCP with first-stream classification for consistent policy enforcement (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Enforce granular application QoS for HTTP/2 traffic using differentiated services code point (DSCP) based on first-stream classifications. The device applies the application classification of the first-stream to the entire HTTP/2 parent session and all subsequent streams. Packets in the session inherit a uniform AppQoS behavior derived from that initial classification. Using first-stream classification ensures that traffic is classified and prioritized by application type when multiple streams use different applications.

This approach integrates with traditional and unified policy frameworks, simplifying deployment without requiring additional configuration.

[See Application QoS.]

HTTP/2 inspection for IDP (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4120, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Intrusion detection and prevention (IDP) now inspects HTTP/2 traffic to enhance security on SRX Series Firewalls. IDP identifies HTTP/2 application traffic and applies relevant attack signatures during traffic inspection. This capability increases visibility into modern web application traffic and improves threat detection over HTTP/2 connections.

IDP inspection secures HTTP/2 traffic within connection sessions and stream sessions. It detects attacks and anomalies at the protocol level.

You can create custom signatures using HTTP/2 contexts such as frame length, frame type, flags, and stream ID.

[See Support for HTTP/2 Traffic Inspection.]

HTTP/2 inspection for SSL proxy (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSL proxy adds Application-Layer Protocol Negotiation (ALPN)-based HTTP/2 inspection with policy controls for TLS 1.2 and TLS 1.3. During the TLS handshake, the proxy negotiates the HTTP/2 identifier h2 with compliant peers and falls back to HTTP/1.1 when HTTP/2 is disabled or not allowed.

HTTP/2 is disabled by default. You can enable or disable it globally by using the *set services ssl proxy global-config http2 on|off* command. *D*isable HTTP/2 selectively per policy using the *set security policies from-zone <from-zone> to-zone <to-zone> policy <policy-name> then permit application-services ssl-proxy disable-http2* command.

The SSL proxy enforces protocols by stripping HTTP/3 ALPN, blocks prohibited TLS 1.2 cipher suites, rejects renegotiation, and terminates TLS 1.3 post handshake authentication.

[See SSL Proxy HTTP/2 Inspection with ALPN (TLS 1.2 and TLS 1.3).]

HTTP/2 protocol support (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—HTTP/2 protocol is a modern version of the HTTP protocol designed to improve application performance and efficiency by enabling multiple transactions over a single connection. Each connection logically divides into multiple streams, enabling your device to inspect, permit, or deny transactions independently while maintaining a single connection.

You can safely enable applications running over HTTP/2 without any additional configuration on your device.

[See HTTP/2 Protocol.]

Support for cleartext HTTP/2 inspection (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Inspect unencrypted HTTP/2 traffic with cleartext HTTP/2 inspection. The firewall identifies HTTP/2 connections early, manages them as single TCP connections, and tracks each HTTP/2 stream independently. The system performs Layer 7 inspection and enforces policies on each stream, providing granular visibility and control over individual transactions within the same connection.

To configure cleartext inspection for HTTP/2, use the CLI commands: `set security application-services http2 plain-text` and `set services ssl proxy global config http2 on`. Verify your configuration with `show security application-services status` and `show security application-services monitoring`.

[See HTTP/2 Inspection: Support for Cleartext HTTP/2 Traffic.]

Support for HTTP/2 traffic in Juniper ATP Cloud (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)—Juniper ATP Cloud supports HTTP/2 traffic, giving you clear, stream-level visibility into HTTP/2

communications. It helps you detect threats accurately, inspect traffic deeply, and apply security policies consistently, without slowing down network performance. This support works seamlessly with Advanced Anti-Malware (AAMW), antivirus, Security Intelligence (SecIntel), and Security Metadata Streaming (SMS).

[See show services advanced-anti-malware statistics, show services anti-virus statistics, and show services security-metadata-streaming.]

Support for HTTP/2 in web filtering and content filtering (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—The HTTP/2 protocol supports Content Security features such as web filtering and content filtering. These features bypass the HTTP/2 parent session and process child sessions directly. The system applies actions based on the existing configuration for specific sessions, ensuring accurate filtering and enhanced security.

[See Content Security Overview and HTTP/2 Protocol.]

Content Security

Deprecation of performance-mode and disable-dynapp-profile-selection statements (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We've deprecated the `performance-mode` and `disable-dynapp-profile-selection` statements from the [`edit security utm default-configuration web-filtering`] hierarchy level.

[See web-filtering.]

Support for dynamic file type signatures in content filtering (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use dynamic file type signatures to efficiently manage file type signature data. These signatures are stored in a structured JSON format on a cloud-based server, allowing you to download, install, and update them without upgrading the existing Junos OS on your device. Before this release, you had to wait for the next Junos OS upgrade to update the file type signature data. Use CLI commands to easily manage tasks such as downloading, installing, and rolling back file type signatures.

[See request security utm content-filtering file-signatures download, request security utm content-filtering file-signatures download status, request security utm content-filtering file-signatures install, request security utm content-filtering file-signatures install status, request security utm content-filtering file-signatures local-update, request security utm content-filtering file-signatures rollback, request security utm content-filtering file-signatures rollback status, request security utm content-filtering file-signatures uninstall, request security utm content-filtering file-signatures uninstall status, and show security utm content-filtering file-signatures status.]

Flow-Based and Packet-Based Forwarding

Optimization of TCP Proxy (cSRX, SRX1600, SRX2300, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We can use SRX Series Firewall as a TCP Proxy. This configuration ensures consistent proxying and prevents implicit services from disrupting sessions. It enhances session reliability and security. To configure security policy with default (global TCP Profile), use the `set security policies from-zone trust to-zone untrust policy p1 then permit tcp-proxy` command.

To configure the TCP proxy profile, use the `set security flow tcp-proxy tcp-profile` command. This command helps in configuring the required parameters.

To verify the configuration, use the`show security flow tcp-proxy profiles` command.

[See TCP Sessions.]

Junos XML API and Scripting

Event policy archive sites support specifying a routing instance (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When you configure an archive site for an event policy destination, you can specify the routing instance that the device uses to connect to the site. To configure the routing instance, include the `routing-instance` option at the `[edit event-options destinations` *`destination-name`* `archive-sites` *`url`*`]` hierarchy level.

[See Event Policy File Archiving and archive-sites (Event Policy).]

Network Address Translation (NAT)

Deterministic NAT with interim logging (cSRX, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—You can now configure syslog intervals for deterministic NAPT. The configurable syslog interval reduces logging requirements by ensuring consistent IP address and port mappings.

Use the `set security nat source pool <pool-name> port deterministic det-nat-configuration-log-interval` command to define intervals from 1800 through 86400 seconds.

[See Deterministic NAT and det-nat-configuration-log-interval.]

Dynamic scaling for persistent NAT capacity (cSRX, SRX1600, SRX2300, SRX4120, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—You can enhance persistent NAT capacity by dynamically scaling beyond the fixed, code-defined limit. Dynamic scaling in persistent NAT enables growth by up to a configurable number of additional entries, depending on the configuration and available runtime resources, in addition to the minimum baseline static bindings allowed on the device.

You must configure the `set security nat source persistent-nat increased-scale-limit <value>` command to increase the persistent NAT capacity.

[See Persistent NAT and NAT64, show security nat source persistent-nat-table, and increased-scale-limit.]

Platform and Infrastructure

Data Plane Development Kit library upgrade (vSRX3.0)—You can now use the DPDK 25.11 version to build vSRX3.0 images.

[See Requirements for vSRX Virtual Firewall on VMware.]

Intel NIC PF PCI passthrough support (vSRX3.0)—Physical Function (PF) PCI passthrough assigns an entire physical NIC port directly to a vSRX3.0 virtual machine (VM) on kernel-based virtual machine (KVM) and VMware ESXi hypervisors. With PF passthrough, the vSRX3.0 data plane accesses the NIC hardware directly, bypassing the hypervisor's virtual switching stack. VMs with PF passthrough NICs do not support live migration. NICs assigned through PF passthrough are dedicated to a single VM and cannot be shared.

This deployment model changes performance characteristics and operational behavior compared to paravirtualized NICs and SR-IOV virtual function (VF) passthrough. Use PF passthrough only with qualified combinations of hypervisor, NIC, and driver.

[See Requirements for vSRX Virtual Firewall on KVM and Requirements for vSRX Virtual Firewall on VMware.]

User Access and Authentication

Sequential DNS fallback for outbound SSH (cSRX, EX3400, EX3400-VC, EX4000-12MP, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48P, EX4000-48T, EX4000-8P, 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, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020,

MX10003, MX10004, MX10008, MX10016, SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX 3.0)—When a configured hostname resolves to multiple IP addresses via DNS, Junos OS automatically tries each address in sequence, regardless of IP family, until a connection succeeds or all addresses have been attempted. The `outbound-ssh` retry mechanism activates after all resolved addresses have been tried.

[See outbound-ssh.]

VPNs

Configurable PRF in IKEv2 with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—You can configure Internet Key Exchange version 2 (IKEv2) proposals with a specific pseudorandom function (PRF) algorithm independent of the authentication algorithm. Use the `prf-algorithm` option to configure the following algorithms at the [`edit security ike proposal` *`proposal-name`*] hierarchy level in the IPsec VPN service with the iked process:

- `prf-hmac-sha-256`
- `prf-hmac-sha-384`
- `prf-hmac-sha-512`
- `prf-hmac-sha1`

You cannot combine a PRF with Authenticated Encryption with Associated Data (AEAD) cipher suites such as AES-GCM or ChaCha20-Poly1305 algorithm.

[See prf-algorithm and proposal (Security IKE).]

Deprecation of certificates as authentication method in IKE (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—We've deprecated support for the `certificates` option at the [`edit security ike proposal` *`proposal-name`* `authentication-method`] hierarchy level. However, you can continue to use the other authentication methods at this hierarchy level.

[See proposal (Security IKE).]

Multiple local certificates in IKE with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—You can specify multiple local certificates in an IKE policy on a device that runs the iked process for the IPsec VPN service. During IKE negotiation, the device selects the local certificate whose certificate chain matches the peer's certificate request. The device evaluates the certificates in their listed order. Use `set security ike policy` *`policy-name`* `certificate local-certificate [`*`cert1 cert2`*`]` to configure multiple local certificates.

[See certificate and IKE Policy with Local Certificate.]

RSA-PSS signature algorithm in IKEv2 with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Specify the RSA signature scheme for IKEv2 authentication with an RSA local certificate as per RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2*. Configure the following algorithms using the `digital-signature-scheme` option at the [`edit security ike proposal` *`proposal-name`*] hierarchy level:

- `rsassa-pkcs1-v1-5` for RSA signature with the PKCS1 v1.5 scheme
- `rsassa-pss` for RSA signature with Probabilistic Signature Scheme (PSS) as the default signature scheme for RSA certificate

Do not configure `digital-signature-scheme` without setting `digital-signature` as the authentication method.

[See digital-signature-scheme, proposal (Security IKE), and show security ike security-associations.]

Support for trusted CA group for certificate validation in IKE with iked process (cSRX, MX240, MX301, MX304, MX480, MX960, MX10008, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Junos OS uses the trusted certification authority (CA) group in an IKE policy to validate peer certificates during IKE authentication. You can use the feature on a device that runs the iked process for the IPsec VPN service. Configure the trusted CA group to specify which CAs the device trusts for incoming peer certificate validation. Use `set security ike policy` *`policy-name`* `certificate trusted-ca trusted-ca-group` *`trusted-ca-group-name`* to reference the trusted CA group in an IKE policy for certificate validation.

[See certificate and IKE Policy with Trusted CA Group.]

Additional Features

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

Support for Quantum Buffer in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100,

SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—Use Juniper Networks Quantum Buffer for JSSH to enhance SSH management and maintain cryptographic agility.

[See The Quantum Buffer and moduli.]

Support for Shor-resistant and other default key exchange algorithms in SSH (ACX5448, ACX5448-D, ACX5448-M, EX3400, EX3400-VC, EX4000-12MP, EX4000-12MUP, EX4000-12MUP-TAA, EX4000-12P, EX4000-12T, EX4000-24MP, EX4000-24MUP, EX4000-24MUP-TAA, EX4000-24P, EX4000-24T, EX4000-48MP, EX4000-48MUP, EX4000-48MUP-TAA, EX4000-48P, EX4000-48T, EX4000-8P, 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-12T, EX4100-H-24F, EX4100-H-24F-DC, EX4100-H-24MP, EX4100-H-24MP-DC, EX4400-24MP, EX4400-24P, EX4400-24T, EX4400-24X, EX4400-48F, EX4400-48MP, EX4400-48MXP, EX4400-48P, EX4400-48T, EX4400-48XP, EX4650, EX4650-48Y-VC, EX9204, EX9208, EX9214, MX204, MX240, MX301, MX304, MX480, MX960, MX2010, MX2020, MX10003, MX10004, MX10008, MX10016, NFX150, NFX250, NFX350, QFX5100VC, QFX5110, QFX5110-VC, QFX5110-VCF, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-48YM, QFX5200, QFX5210, QFX10002-60C, SRX1500, SRX1600, SRX2300, SRX4100, SRX4120, SRX4200, SRX4300, SRX4600, SRX4700, SRX5400, SRX5600, SRX5800, and vSRX3.0)—SSH supports the Shor's algorithm-resistant hybrid Streamlined NTRU Prime 761 and X25519 key exchange algorithm. And, by default, SSH supports the `dh-group16-sha512` and `dh-group18-sha512` Diffie-Hellman (DH) group key exchange algorithms.

[See key-exchange.]

What's Changed

IN THIS SECTION

Learn about what's changed in this release for vSRX.

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

EVPN

- **Pseudowire status display update in show evpn vpws-instance output**—For EVPN-VPWS pseudowire interfaces configured with `protocols ccc`, the Pseudowire field in the `show evpn vpws-instance` command output now displays the correct corresponding pseudowire status value Not Present when the local logical interface for the pseudowire connection is not present. Previously, this field still displayed the value CCC-Up even when the Status output field showed the local interface was not present.

  [See show evpn vpws-instance.]

- **Preference-based DF election designated-forwarder-preference-xxxx options set at the global [edit protocols evpn] hierarchy level apply only to the default switch instance**—You can configure the `designated-forwarder-preference-highest` or `designated-forwarder-preference-least` option at the `[edit protocols evpn]` hierarchy level for EVPN preference-based designated forwarder (DF) election. You can alternatively set these options at the `[edit routing-instances name protocols evpn]` hierarchy level for a particular routing instance. Setting either of these options at the `[edit protocols evpn]` hierarchy level applies to EVPN multihoming only in the default switch instance. The configuration doesn't apply globally to all defined EVPN routing instances on the device. These options are available at the `[edit protocols evpn]` hierarchy level only on platforms that support a default switch instance. To enable these options for EVPN routing instances other than a default switch instance, you must configure them in each routing instance at the `[edit routing-instances name protocols evpn]` hierarchy level.

  [See Preference-Based DF Election.]

- **XML output tags changed for `show l2-learning evpn mcnh-info` command**—We've changed the XML output for this command. The XML `<l2ald-mcnh-entry-list-headers>` element level under the `<l2ald-mcnh-entry>` element has been removed. The child output tags that were under the `<l2ald-mcnh-entry-list-headers>` element are now direct child tags of the `<l2ald-mcnh-entry-list>` element.

General Routing

- **Enhancements to port-extender configuration statement**—You can see the supported platforms when configuring a satellite device within a NextGen Port Extender topology. At the `[edit services port-extender satellite <name> device-model]` hierarchy, you can use the '?' to see a list of supported platforms.

We have also added the `skip-lo0-config` statement under [`edit service port-extender jnu`] hierarchy in Junos OS Release 25.4R2 and later. When this knob is configured, the commit script skips updating the lo0 configuration. We have added `cp-mtu` under the [`edit service port-extender satellite <name>`] hierarchy in Junos OS Release 25.4R2 and later, allowing you to configure an MTU value other than the default MTU.

[See port-extender (Services).]

- **Router Advertisement Behavior Adjustment**—The system now sends a multicast Router Advertisement (RA) in response to IPv6 Router Solicitations (RS) with an unspecified source address (::) under the solicit-router-advertisement-unicast configuration. This adjustment ensures compliance with RFC 4861 by providing hosts without IPv6 addresses with the necessary router information and prefixes for auto-configuration. When an RS packet is received from an unspecified source (::), the router will send an RA to the all-nodes multicast address `ff02::1`. The RA source is the router's link-local address, and the Source Link-Layer Address (SLLA) option is omitted. For RS packets with valid unicast IPv6 sources, the router continues sending unicast RAs as before. No configuration changes are required, and existing setups will automatically benefit from this update.

- You can select RSA or ED25519 encrypted private and public key pairs (.pem) when provisioning instances on AWS and GCP cloud platforms. This option meets security policies and cryptographic preferences and provides secure, key-based access. Generate or import the desired key pair and attach the key pair during instance creation to authorize administrative access.

- **Improved route opaque client identification**—The opaque data client display field has been enhanced so that opaque clients that do not have a valid client ID are now labeled with their TLV type value instead of Unknown (0). This change improves the clarity and debuggability of route opaque information, especially when multiple opaque components are present on a route. Existing opaques that use a client ID continue to display their client name or ID. Opaque data and Opaque data client can be viewed by issuing the `show route extensive expanded-nh` command.

  [See show route extensive <route> expanded-nh.]

- **New fields in the `show interfaces interface`**—The `show interfaces <interface-name>` command output introduces two new fields to identify the custom media code and the host code. Here's a sample output

```
user@host> show interfaces interface-name
Physical interface: interface-name, Enabled, Physical link is Up
  Interface index: 1249, SNMP ifIndex: 533
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 400Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
```

```
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled, Auto-negotiation:
Disabled, Media type: Fiber
  Wavelength     : 1550.12 nm, Frequency: 193.400 THz
  Optic-loopback : Disabled , Optic-loopbacktype : nil
  Media Code     : ZR400-OFEC-16QAM-HB(54)              - 54 is the media id, "ZR400-
OFEC-16QAM-HB" is the description of this media.
  Host Code      : 400GAUI-8 C2M (Annex 120E)(17)       -17 is the host id, "400GAUI-8 C2M
(Annex 120E)" is the description of this host.
  Device flags   : Present Running
```

- The priority order for VSTP configurations has changed when the same VLAN is defined across multiple configuration levels simultaneously (vlan-id, vlan-group <group-name>, and vlan all). In releases earlier than Junos OS Release 26.2R1 and Junos OS Evolved Release 26.2R1, vlan-group configurations used to take precedence over vlan all. The updated behavior ensures that a VLAN group takes precedence over the vlan all statement.

- [New output field added to the `show system switchover` command—The `show system switchover` command introduces a new output field for fabric readiness status:. This change is applicable to Evo PTX products. user@host> system switchover Graceful switchover: On Configuration database: Ready Object database: Ready Applications' ready state: Ready Commit in progress: False Switchover Status: Ready Fabric ready status: Ready >>>>>>>>>>> New tag for fabric readiness status >>>>>>>>>>>>>

- **Per-fan entries for fan trays in SNMP MIB tables (PTX10008)**—You can monitor individual fans in a fan tray through distinct entries in SNMP MIB tables. This enhancement addresses the earlier limitation that SNMP MIB does not list individual fans.

- **Suppress periodic export of IPFIX records**—When this option is enabled, IPFIX records are exported only when:

  - The software has learned a new flow.

  - A flow has aged-out.

  Suppressing periodic export allows more controlled and event-driven export of IPFIX records instead of continuous periodic transmission. To suppress the periodic export of IPFIX records, configure the `suppress-periodic-export` statement at the `[edit services inline-monitoring template <i>template-name</i>]` hierarchy level.

  [See template (Inline Monitoring).]

- **Member Link Limit on Adaptive Aggregated Ethernet (PTX Series)**—A maximum of 64 member links are configurable on an aggregated Ethernet (AE) bundle with adaptive load balancing. Configuring more than 64 member links will result in a commit failure.

[See Understanding Aggregated Ethernet Load Balancing.]

- Write cache disabled alarms no longer raised—The "Disk 2 Write Cache disabled" alarms were reported in the chassis alarms due to a missing SSD model in Junos OS supported device table. The support is now added, and alarms are no longer reported.

Layer 2 Ethernet Services

- **New file path for customer binding information**—The Dynamic Host Configuration Protocol (DHCP) local server or relay agent now stores customer binding information to a file in **/var/db**. Previously, the file was stored in **/var/preserve**, which was not secured. To restore the customer binding information after an upgrade or reboot, choose one of the following approaches: -Before upgrading, copy the file from **/var/preserve** to **/var/db**. -After upgrading, copy the file from **/var/preserve** to **/var/db** and run the `request dhcp restore-persistent-clients from-file` command.

  [See Configuring DHCP Local Server to Preserve Subscriber Binding Information.]

MPLS

- **Mandatory MPLS IPv6 Traffic Engineering Configuration**—New configuration requirements have been introduced to support the export and import of IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses within the traffic engineering database (TED). The new configuration requirements are as follows:

  - The `set protocols mpls traffic-engineering database export ipv6` command is now visible and mandatory to export IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from lsdist to TED.

  - The `set protocols mpls traffic-engineering database import ipv6` command is mandatory to import IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses from TED to lsdist.

  - The `set protocols isis traffic-engineering ipv6` command is mandatory to allow IS-IS-learned IPv6 prefixes, SRv6 END and END.X SIDs, and node or link IPv6 addresses to populate TED.

  These configurations ensure complete IPv6 coverage across TED, IS-IS, and BGP-LS workflows, including proper propagation of IPv6 router identifiers from BGP-LS node attribute TLVs into TED.

  [See export (MPLS Traffic engineering database), import (MPLS Traffic Engineering Database), and traffic-engineering (Protocols IS-IS).]

Network Management and Monitoring

- **Improved Syslog User Identification (SRX4300)**—Syslogs generated from shell commands now logs the remote username instead of the local template username when you use the RADIUS or TACACS + authentication.

- **Improved ZTP SNMP configuration handling (Junos)**—We have reordered the SNMP configuration commands to optimize the Zero Touch Provisioning (ZTP) process. The `set snmp engine-id local` statement now processes before `set snmp v3 usm local-engine user`, ensuring SNMPv3 USM authentication and privacy keys generation using the configured local engine ID instead of an auto-derived default. This prevents SNMPv3 query authentication failures from engine-ID resets.

  [See Configure Engine ID on SNMPv3 and Create SNMPv3 Users.]

Platform and Infrastructure

- **ARP proxy default reply option**—We have introduced the `proxy-arp-reply-for-default` option under the `set system arp` hierarchy to enable the device to reply to proxy ARP requests for default routes.

Routing Protocols

- **Enhanced BGP Negotiated Add-Path Capabilities output**—The `show bgp neighbor`command output has been enhanced to display whether add-path `send` and `receive` capabilities are negotiated or not for each NLRI. These enhancements provide improved visibility into the BGP session's add-path capabilities, aiding in correctly troubleshooting BGP configurations.

  [See show bgp neighbor.]

- **Handling of zero Link Bandwidth Extended Community values for WECMP**—When zero LBW values are received on a route, the balance value is set to zero only for that specific path within WECMP, rather than falling back to ECMP. This change aligns with the latest IETF recommendations for weighted load-balancing.

  [See BGP Link-Bandwidth Community.]

- **BGP Route Import Policy Update**—The handling of BGP route import policies has been updated where the peer import policy is applied only once in the primary RIB, while the secondary RIB applies only the rib-group import policy, if any. This change ensures the accurate handling of AS-path and additive attributes such as MED, and local-preference therefore improving policy consistency and routing accuracy. Previously, the primary import policy was re-applied in the secondary RIB through rib-groups, causing duplicate policy effects.

- **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.]

- Supports static route configurations with /32 host addresses—Static Route Configuration with /32 host addresses allows configuring static routes using an interface as the next-hop, even if the interface has a /32 netmask. This change is significant for network administrators managing static

routes in environments where interfaces are often assigned /32 addresses. Routes configured in this manner will now commit successfully, removing previous limitations. Example for CLI commands of static routes that uses an interface addresses as next-hop: set routing-options static route 45.63.100.176/32 next-hop et-0/0/32.0 set interfaces et-0/0/32 unit 0 family inet address 169.254.255.32/32

Unified Threat Management (UTM)

- **Support for IPv6 in Enhanced Web filtering (EWF) (SRX Series)**—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.

- Dynamic application policies with shared content filtering Content Security policy and ALG data application handling (SRX Series Firewalls and vSRX3.0)—If you configure dynamic applications in the policy match conditions, you do not need to explicitly specify the application (L4 port), because the system determines it automatically. However, for ALG (Application Layer Gateway) data applications, such as junos:FTP-DATA, the system cannot automatically identify the L4 port. In these cases, you must either set the match condition to application any or use custom L4 ports for filtering. If you require L7 services on ALG data sessions or want to match dynamic applications, the recommended approach is to configure the match condition as application any.

  [See Content Security Overview.]

- 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.]

- **Enhanced anti-virus and web-filtering operational command outputs (SRX Series Firewall and vSRX)**—We've enhanced the `show security utm web-filtering category status` and `show security utm anti-virus status` outputs. The enhanced web-filtering output displays error message for the failed category-package file download. When you enable Avira engine, the enhanced anti-virus output displays error messages for the failed Avira pattern-updater download.

  Proxy profile settings update (host, port, username, password) causes download issues or connectivity problems, which are accurately reflected in the anti-virus and web-filtering operational command outputs.

  [See show security utm web-filtering category status and show security utm anti-virus status.]

User Interface and Configuration

- **Changes to the `show system data-models` command and corresponding RPC**—We've made the following changes to the `show system data-models genstate` command and RPC to better represent the data and command options for genstate models:

  - Renamed the command option from `xpath-cli-command-mapping` to `genstate-path-cli-command-mapping`
  - Changed the RPC to rename the request tag from `<get-genstate-xpath-cli-command-mapping>` to `<get-genstate-path-cli-command-mapping>` and rename the `<xpath>` option to `<genstate-path>`
  - Changed the XML output to rename the root element from `<genstate-xpath-cli-command-mapping-information>` to `<genstate-path-cli-command-mapping-information>` and rename the `<xpath>` element to `<genstate-path>`

  [See show system data-models.]

- **Introduction of remove operation**—A new CLI operation `remove` has been introduced to handle the deletion of configurations more efficiently. This operation prevents <userinput>statement not found</userinput> warnings when the target configuration does not exist, eliminating potential issues with RPC responses and avoiding buffer overflows during script operations, making configuration management more robust.

- **CLI options and RPC formats to retrieve `genstate` information and XPath expressions for operational commands**—A gNMI telemetry collector can subscribe to `genstate` resource paths to query for the corresponding state data on a device. We've added CLI pipe filter functions and RPC formats that enable you to discover available `genstate` resource paths and the associated `genstate` YANG module for a given show command. You can also retrieve the XPath expression for any node in command output. The options include:

  - `| display genstate subscription-paths`—Display all resource paths available for subscription in the `genstate` YANG module associated with the given command.
  - `| display genstate yang-module-name`—Display the name of the `genstate` YANG module that includes the data model and resource paths for the given command.
  - `| display xml genstate-path`—Display the command output in XML and include the `junos:genstate-path` attribute for every emitted tag that is available for subscription. The attribute value is the `genstate` subscription path. For RPCs, use the `format='xml-genstate-path'` attribute in the request tag.
  - `| display xml rpc response-paths`—List all possible response paths under the equivalent RPC's top-level response tag.
  - `| display xml xpath`—Display the command output in XML and include the `junos:data-xpath` attribute for each tag in the output. The attribute value is the XPath expression that selects that specific node. For RPCs, use the `format='xml-xpath'` attribute in the request tag.

[See | (pipe) and Junos Genstate YANG Data Models.]

VPNs

- **Updated show route extensive output field for MVPN c-Multicast Routes**—The `show route extensive` output for MVPN c-mcast routes `Inactive reason` is updated to `MVPN Active preferred`. Earlier it displayed the reason as `Not Best in its group - Active preferred`.

  [See show route extensive.]

- **Removal of weak algorithms in IPsec VPN with iked process (MX301, MX304, MX10004, MX10008, MX Series with MX-SPC3, NFX250, NFX350, SRX Series, and vSRX3.0)**—The following algorithms are no longer available as CLI options:

  - `des-cbc` and `3des-cbc` as encryption algorithms in IKE proposal and IPsec proposal at the `[edit security ike proposal encryption-algorithm]` and `[edit security ipsec proposal encryption-algorithm]` hierarchy levels respectively.
  - `md5` and `sha1` as authentication algorithms in IKE proposal at the `[edit security ike proposal authentication-algorithm]` hierarchy level.
  - `group1`, `group2`, and `group5` as DH groups in IKE proposal and Perfect Forward Secrecy (PFS) groups at the `[edit security ike proposal dh-group]` and `[edit security ipsec policy perfect-forward-secrecy keys]` hierarchy levels respectively.
  - `hmac-md5-96`, `hmac-sha1-96`, and `hmac-sha-256-96` as authentication algorithms in IPsec proposal at the `[edit security ipsec proposal authentication-algorithm]` hierarchy level.
  - `basic`, `compatible`, and `standard` as IKE proposal sets and IPsec proposal sets at the `[edit security ike policy proposal-set]` and `[edit security ipsec policy proposal-set]` hierarchy levels respectively.

  For MNHA HA link encryption and standalone VPNs, using `hmac-sha1-96` as the IPsec authentication algorithm at the `[edit security ipsec proposal proposal-name authentication-algorithm]` hierarchy level results in a commit error. Configure stronger algorithms to enhance IPsec VPN security.

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

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for vSRX.

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

J-Web

- J-Web becomes unreachable when a local self-signed certificate with weak digital signature is assigned to the HTTPS PKI configuration. PR1955001

Open Issues

IN THIS SECTION

Learn about open issues in this release for vSRX.

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

Intrusion Detection and Prevention (IDP)

- IDP: On-box packet capture fails when packet-log host or port is not reachable with SSL Initiation Profile For On-box packet capture, ssl initiation profile configuration is NOT required. This issue will not be seen if we do not configure ssl initiation profile. PR1921354

Resolved Issues

IN THIS SECTION

Learn about resolved issues in this release for vSRX.

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

General Routing

- On all SRX platforms, the firewall cps (connection per second) drop is seen. PR1935299
- On vSRX platforms, enabling SSL (Secure Sockets Layer) and AppSec (Application Security) may result in an approximate 6-8% performance degradation. This behavior is linked to the HTTP2 (HyperText Transfer Protocol) module default configuration. When the HTTP2 is enabled, additional processing overhead may lead to the observed performance drop. PR1912321
- When firewall authentication is used in policy with client sending the traffic over HTTP/2, the firewall-auth plugin of SRX temporarily downgrades the connection to HTTP/1.1 to complete user authentication. After successful authentication, the session resumes in HTTP/2. For firewall authentication to work correctly, the client must allow default HTTP protocol negotiation options. Config for Firewall-auth in policy : set security policies from-zone trust to-zone untrust policy policyteampolicy then permit firewall-authentication <..> PR1934371
- Starting in 25.4R2, customers can use ed25519 ssh keys from either the web console or CLI of AWS and GCP as well as rsa ssh keys when selecting an SSH key to be able to log into the system with at launch. Configuring the keys through cloud-init config was already supported. PR1922165
- Upon booting up a Junos vSRX platform following an upgrade to Junos 24.4R2-S3 and above, one or more interfaces may not be seen within the interfaces list from running a "show interfaces terse" for aggregate ethernet (AE) interfaces. Due to this it causes AE bundling issue and interfaces don't come up leading to traffic loss. PR1945668
- On Junos OS and Junos OS Evolved platforms, an issue has been identified where post-dated licenses do not automatically activate when their configured start date is reached. This issue occurs

when a new license is installed with a future activation date. Once the start date arrives, the license remains inactive instead of being automatically enabled as expected. PR1950936

Platform and Infrastructure

- On Junos and Junos OS Evolved platforms, the output format of the fetch command changed after upgrading to a newer upstream version of the utility. The earlier format displayed an explicit completion indicator (for example, 100%). The upgraded version removed this indicator. PR1924159

Unified Threat Management (UTM)

- On Junos OS SRX platforms, configuring cache preload in web filtering causes high memory utilization in the UTM (Unified Threat Management) pool, which results in traffic loss. PR1927484

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

Basic Procedure for Upgrading to Release 26.2R1

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

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

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

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

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: 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-x86-32-26.2R1.9-signed.tgz
      ```

    - For 64-bit Routing Engine version:

      ```
      user@host> request system software add no-validate reboot source/junos-install-mx-x86-64-26.2R1.9-signed.tgz
      ```

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

    - For 32-bit Routing Engine version:

      ```
      user@host> request system software add no-validate reboot source/junos-install-mx-x86-32-26.2R1.x-limited.tgz
      ```

    - For 64-bit Routing Engine version:

      ```
      user@host> request system software add no-validate reboot source/junos-install-mx-x86-64-26.2R1.9-limited.tgz
      ```

    Replace source with one of the following values:

    - */pathname*—For a software package that is installed from a local directory on the router.
    - For software packages that are downloaded and installed from a remote location:
      - ftp:// */pathname*
      - http:// */pathname*
      - scp:// */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.

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: 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-x86-32-26.2R1.9-signed.tgz
      ```

    - For 64-bit Routing Engine version:

      ```
      user@host> request system software add no-validate reboot source/junos-install-mx-x86-64-26.2R1.9-signed.tgz
      ```

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

- For 32-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot source/junos-install-mx-x86-32-26.2R1.x-limited.tgz
  ```

- For 64-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot source/junos-install-mx-x86-64-26.2R1.9-limited.tgz
  ```

  Replace source with one of the following values:

- */pathname*—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - ftp:// */pathname*
  - http:// */pathname*
  - scp:// */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.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

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

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

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

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

For more information about EOL and EEOL releases, see Junos OS Dates & Milestones - Juniper Networks.

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

We have two types of releases, EOL and EEOL:

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

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

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

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

For more information about EOL and EEOL releases, see Junos OS Dates & Milestones - Juniper Networks.

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

Upgrading a Router with Redundant Routing Engines

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

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

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

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 26.2R1

To downgrade from Release 26.2R1 to another supported release, follow the procedure for upgrading, but replace the 26.2R1 jinstall package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the Installation and Upgrade Guide.

To downgrade from Release 26.2R1 to another supported release, follow the procedure for upgrading, but replace the 26.2R1 jinstall package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the Installation and Upgrade Guide.

Documentation Updates

There are no documentation updates for this release.

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.

Feature Explorer

- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved.

  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.

  Compliance Advisor.

Requesting Technical Support

IN THIS SECTION

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit Product Warranty Policy - Support - Juniper Networks.
- 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: Support
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: Documentation
- Find solutions and answer questions using our Knowledge Base: CEC Juniper Community
- Download the latest versions of software and review release notes: Support
- Search technical bulletins for relevant hardware and software notifications: CEC Juniper Community
- Join and participate in the Juniper Networks Community Forum: Elevate - Elevate Community

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

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: Support
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: Documentation
- Find solutions and answer questions using our Knowledge Base: CEC Juniper Community
- Download the latest versions of software and review release notes: Support
- Search technical bulletins for relevant hardware and software notifications: CEC Juniper Community
- Join and participate in the Juniper Networks Community Forum: Elevate - Elevate Community

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 Contact - Support - Juniper Networks.

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 Contact - Support - Juniper Networks.

Revision History

29-Jun-26 - Revision 1, Junos OS Release 26.2R1.