Release Notes

Published
2026-01-08

Junos OS Evolved Release 23.1R1

Introduction

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 23.1R1.

For more information on this release of Junos OS Evolved, see Introducing Junos OS Evolved.

Table of Contents

Junos OS Evolved Release Notes for ACX7024, ACX7100-32C, ACX7100-48L, and ACX7509 Devices

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 23.1R1 for ACX7024, ACX7100-32C, ACX7100-48L, and ACX7509 devices. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

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

To view features supported on the ACX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Evolved Release 23.1R1, click the Group by Release link. You can collapse and expand the list as needed.

- ACX7024
- ACX7100-32C
- ACX7100-48L
- ACX7509

The following sections highlight the key features in this release.

EVPN

- **Automatically derived ESI support on EVPN-MPLS (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, you can configure multihomed devices in an EVPN-MPLS network to automatically generate the ESI from the:

  - System ID and administrative key on the remote customer edge (CE) device (partner).
  - Locally configured MAC and local discriminator values.

  [See Other Methods to Auto-Derive the ESI.]

- **Support for DHCP-relay deployment of EVPN over MPLS that includes ERB (ACX Series)**

  —Starting in Junos OS Evolved Release 23.1R1, you can connect DHCP clients to EVPN spine provider edges (PEs) directly. These spine PEs also perform DHCP relay functions, and the routers support spine functionality for Integrated Routing and Bridging (IRB).

  We support the following functionalities:

  - EVPN over MPLS Ethernet-LAN
  - DHCPv4 and DHCPv6 relay options
  - Stateless forward-only mode for DHCP relay over VPN
  - Anycast IP address with IRB for a relay-source
  - Client VRFs only

  [See DHCP Relay Agent in EVPN-MPLS Network.]

- **Determine IRB interface state changes based on local and remote connectivity states in EVPN fabrics (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700)**—Starting in Junos OS Evolved Release 23.1R1, the provider edge (PE) devices in an EVPN fabric consider the following factors when determining the state (up or down) of an L3 integrated routing and bridging (IRB) interface. These factors apply to an L3 IRB interface that is associated with a bridge domain or a VLAN in an EVPN instance (EVI).

  - Associated local L2 interface states

    To customize the L2 interface name and other parameters that the device uses to compute the IRB interface state, configure the `interface-state` statement at the `[edit interfaces irb unit n]` hierarchy level.

  - Remote PE device reachability based on the network isolation state of the bridge domain or the EVI

    The device includes the states of the associated EVPN overlay tunnel interfaces in the network isolation state evaluation.

    To define the parameters that determine when an EVI or a bridge domain is in a network isolation state:

    1. Configure the `network-isolation group group-name` statement at the `[edit protocols]` hierarchy level to define a network isolation profile using the available options.
    2. Assign the network isolation group profile to a bridge domain or an EVI using the `network-isolation-profile group network-isolation-group-name` statement at these hierarchy levels:
       - Bridge domain—`[edit bridge-domain bd-name bridge-options]`
       - EVI—`[edit routing-instance instance-name switch-options]`

  [See Determine IRB Interface State Changes from Local and Remote Connectivity States in EVPN Fabrics, interface-state, and network-isolation.]

High Availability

- **Redundancy support enhancements (ACX7509)**—Starting in Junos OS Evolved Release 23.1R1, ACX7509 Series routers have enhanced redundancy support and faster switchover speed. The `show system switchover` and `request chassis routing-engine master switch check` commands have been modified to reflect new switchover behavior. Both commands now display the status of the backup FEB and associated PFEs, including the sync status of the PFEs on the backup FEB.

  [See show system switchover.]

Junos Telemetry Interface

- **Number of configurable BMP monitoring stations increases to a maximum of eight (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, and QFX5220)**—Starting in Junos OS Evolved Release 23.1R1, Junos telemery interface (JTI) delivers initial sync and ON_CHANGE BGP routing information base (also known as routing table) statistics by using remote procedure calls (gRPC) or the gRPC network management interface (gNMI) from a device to an outside collector for a maximum of eight BMP monitoring stations.

- **Logical subinterface sensor support (ACX7100-32C, ACX7100-48L, ACX7509, and ACX7024)**—Starting in Junos OS Evolved Release 23.1R1, Junos telemetry interface (JTI) introduces support for subinterface state and logical interface counters for IPv6 statistics.

  [See Telemetry Sensor Explorer for sensor information.]

- **Platform sensors for chassis components (ACX7509 and PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.1R1, Junos telemetry interface (JTI) supports operational state sensors for chassis backplane, fabric, fan, power supply, and storage components based on the OpenConfig model **openconfig-platform.yang**. Telemetry data for each component is streamed under the path `/components/component[name=`*`component-name`*`]/`.

  [See Telemetry Sensor Explorer.]

Network Management and Monitoring

- **Minimize system reboot systemd console logging (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220)**—Starting in Junos OS Evolved 23.1R1, we support writing all system application start and application stop messages to the console, not just the error messages. During system reboot, by default, the systemd process logs messages for applications to the journal log. By default, application stop and application start status messages are written to the journal log. Only application stop and application start errors are logged to the console. To change this default, configure the `application-status any` statement at the `[edit system syslog console]` hierarchy level.

  [See application-status.]

- **Logging support for routing engine shell and line card shell (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, and QFX5130-32CD)**—Starting in Junos OS Evolved Release 23.1R1, you can log commands executed from the shell when you configure `set system syslog shell`.

  [See syslog (System) and start shell.]

Routing Protocols

- **IS-IS Maximum LSP Size (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 23.1R1, you can configure the maximum LSP size for IS-IS in the range 512 through 9192 bytes to support advertisement of a higher number of prefixes.

  [See max-lsp-size.]

- **Support for blocking the redistribution from a specific protocol into IS-IS (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 23.1R1, you can block the export policy from redistributing the routes from other protocols into IS-IS. You can block the redistribution using the `set protocol isis no-external-export` protocol statement at the [`set protocols isis`] hierarchy level.

  [See no-external-export (Protocols IS-IS).]

- **Support for preventing IS-IS from entering overload state on reaching prefix limit (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 23.1R1, you can prevent IS-IS from entering the overload state even after reaching the set prefix limit. You can configure the `set protocols isis dynamic-overload no-overload-on-prefix-export-limit` statement at the [`set protocols isis`] hierarchy level.

  [See no-overload-on-prefix-export-limit (Protocols IS-IS).]

- **Support for auto-recovery from IS-IS overload state (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 23.1R1, you can set up IS-IS to automatically exit an overload state. The command prevents IS-IS from flushing all the LSP fragments on overload so that when the fragment space is available, IS-IS automatically exits from the overload. You can configure the auto-recovery using the `set protocols isis dynamic-overload auto-recovery` statement at the [`set protocols isis`] hierarchy level.

  [See auto-recovery (Protocols IS-IS).]

- **Support for policy-based ORR (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220)**—Starting in Junos OS Evolved Release 23.1R1, the policy-based ORR helps you select the paths to advertise to achieve your traffic-engineering requirements. You can specify a subset of the paths as candidate paths for path selection. The existing path-selection algorithms select the best paths, and you can modify the attributes of the selected paths or reject the selected paths. The policy-based ORR can work alone or with IGP-based ORR and add-path.

  To enable this feature in BGP peer groups, configure `export <policy>` at the [`protocols bgp group <name> optimal-route-reflection`] hierarchy level.

  [See export (Protocols BGP), optimal-route-reflection, and show bgp group.]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Support for unnumbered interfaces for IS-IS with SPRING TI-LFA (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**— Starting in Junos Evolved OS Release 23.1R1, we support IS-IS over unnumbered interfaces on point-to-point links with SPRING topology-independent loop-free alternate (TI-LFA). You can configure unnumbered interfaces to share the same subnet across multiple interfaces to conserve IPv4 addresses. Note that we do not currently support unnumbered interfaces for IPv6.

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

Additional Features

We have extended support for the following features to the platforms shown in parentheses:

- **CoS support for EVPN-VXLAN** (ACX7024).

  [See CoS Support on EVPN VXLANs .]

- **EVPN-VXLAN support** (ACX7024). We support the following features:

  - VXLAN Layer 2 gateways that provide Layer 2 connectivity between leaf devices separated by an underlay Layer 3 network

    [See Understanding EVPN with VXLAN Data Plane Encapsulation.]

  - Layer 3 Gateway Encapsulation support for Layer 3 over an IPv4 underlay

    [See Understanding VXLANs.]

  - EVPN-VXLAN support for Active/Active multihoming with Ethernet Switch Identifier for Link Aggregation Groups (ESI-LAG)

    [See EVPN VXLAN User Guide - Leaf Device Forwarding.]

  - Storm control support in an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) overlay

    [See Understanding Storm Control.]

  - VLAN-aware bundle , VLAN-based and VLAN bundle MAC-VRF service types

    [See MAC-VRF Routing Instance Type Overview.]

  - Spine role with edge-routed bridging (ERB)

    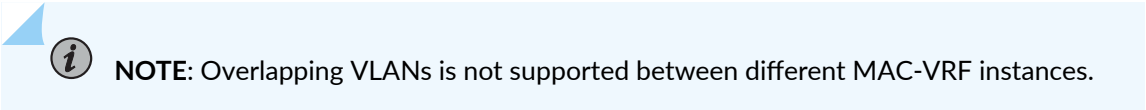

    NOTE: Overlapping VLANs is not supported between different MAC-VRF instances.

  - BPDU protection in EVPN-VXLAN

[See Understanding BPDU Protection for EVPN-VXLAN.]

- Preference-based DF election

  [See Designated Forwarder Election.]

- EVPN BGP policy options support for control plane
- Service provider style interface configuration with VXLAN, including:
  - Encapsulation types ethernet-bridge, extended-vlan-bridge, and flexible-ethernet-services

    [See encapsulation.]

  - Flexible VLAN tagging for single- and dual-tagged logical interfaces on the same port.
  - Q-in-Q

    [See Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network and Understanding Flexible Ethernet Services Support With EVPN-VXLAN.]

- Extensions to enterprise style interface configuration support with VXLAN, including:
  - Trunk interfaces with flexible Ethernet services and flexible VLAN tagging
  - Multiple trunk interfaces on the same port

    [See Understanding Flexible Ethernet Services Support with EVPN-VXLAN.]

- Layer 2 (L2) and Layer 3 (L3) logical interfaces on the same port, with L2 logical interfaces as part of the VXLAN VLAN or tenant VLANs

  [See EVPN User Guide.]

- **Flow-aware transport for pseudowires (FAT) label and entropy label support for Layer 2 circuit and Layer 2 VPN** (ACX7100-32C, ACX7100-48L, and ACX7509 devices).

  FAT label support includes these features:

  - Transmit and receive direction on Layer 2 circuits and Layer 2 VPNs.
  - Layer 2 VPN with multihoming.
  - Layer 2 circuits with redundant pseudowire.
  - Load balancing over LAG.
  - Static flow label for Layer 2 circuits.

  Entropy label support that includes these features:

  - Entropy label indication for LDP and RSVP transport.
  - Load balancing over LAG.

  [See FAT Flow Labels Overview and Configuring Entropy Labels.]

- **Inline active flow monitoring support for IPFIX and v9 export formats** (ACX7509, ACX7100). We support ingress and egress sampling of IPv4 and IPv6 traffic on aggregated Ethernet and IRB interfaces, for both the IPFIX and version 9 export formats. You can configure up to four IPv4 collectors for inline active flow monitoring.

  [See Understand Inline Active Flow Monitoring.]

- **IPv4 address conservation method for hosting providers** (ACX7100). Support for static route with next hop as integrated routing and bridging (IRB) over VLAN and VXLAN to significantly reduce the number of routable IPv4 addresses that you use for your hosting customers.

  [See IPv4 Address Conservation Method for Hosting Providers, Static Route Preferences and Qualified Next Hops, and Configure Static Routes.]

- **MACsec on logical interfaces** (ACX7100-32C).

  [See Media Access Control Security (MACsec) over WAN.]

- **OpenConfig interface configuration support** (ACX7100-32C, ACX7100-48L, and ACX7509).

  [See Mapping OpenConfig Interface Commands to Junos Configuration.]

- **Platform Resiliency Support** (ACX Series Devices). We include chassis synchronization and fault-handling support for components such as:

  - Fan.
  - Clocking.
  - PTP.
  - FPGA.

  With resiliency enabled, the device can detect and log errors, raise alarms, send SNMP traps, provide error indication through LEDs, self-heal, and take components out of service.

  [See show system errors active.]

- **Prevent script execution based on current system memory usage** (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700). You can configure the system memory usage threshold above which the device prevents the execution of certain op, event, SNMP, or Juniper Extension Toolkit (JET) scripts. [See Configure Script Start Options.]

- **Real-time Performance Monitoring (RPM) IPv6 source and target address support** (ACX7024, ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700). You can configure IPv6 source and target addresses for RPM. We've also added support for IPv6 addresses to the SNMP RFC2925a MIB control and results tables. For IPv6 RPM probes, you can enable timestamps only in the Routing Engine.

  [See source-address and target.]

- **Support for BGP Classful Transport (BGP CT)** (ACX7100-32C, ACX7100-48L, ACX7509, and ACX7024 Devices). Support for BGP transport address family or BGP Classful Transport (BGP CT) that includes:

  - Support for service mapping over colored transport tunnels (RSVP, IS-IS flexible algorithm) to transport classes and map service routes over an intended transport class. The transport tunnels can span multiple domains(autonomous systems (ASs) or IGP areas.
  - Network slicing and interoperability between network domains.
  - IPv6 and Segment Routing TrafficEngineered (SR-TE) color-only support.
  - IPv6 and BGP service-routes with a color-only mapping community.
  - Enhanced transport-class configuration to provide precise resolution.

  [See use-transport-class.]

- **SRv6 network programming in IS-IS** (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024). You can configure segment routing in a core IPv6 network without an MPLS data plane.

  To enable SRv6 network programming in an IPv6 domain, use the `source-packet-routing srv6` command.

  To advertise the Segment Routing Header (SRH) locator with a mapped flexible algorithm, use the algorithm at `[edit protocols isis source-packet-routing srv6 locator]` hierarchy level.

  To configure a topology-independent loop-free alternate backup path for SRv6 in an IS-IS network, include the `transit-srh-insert` statement at the `[edit protocols isis source-packet-routing srv6]` hierarchy level.

  [See source-packet-routing srv6 and How to Enable SRv6 Network Programming in IS-IS Networks.]

- **SRv6 support for static SR-TE policy** (ACX7100-32C, ACX7100-48L, ACX7509, and ACX7024 routers). You can configure static segment routing–traffic engineering (SR-TE) tunnels over an SRv6 data plane.

  [See Understanding SR-TE Policy for SRv6 Tunnel.]

- **SRv6 network programming in IS-IS** (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024 with MPC7E, MPC8E and MPC9E line cards). You can configure segment routing in a core IPv6 network without an MPLS data plane.

  To enable SRv6 network programming in an IPv6 domain, use the `source-packet-routing srv6` command. [See source-packet-routing srv6.]

  To advertise the Segment Routing Header (SRH) locator with a mapped flexible algorithm, use the algorithm `[edit protocols isis source-packet-routing srv6 locator]` hierarchy level.

  To configure a topology-independent loop-free alternate backup path for SRv6 in an IS-IS network, include the transit-srh-insert statement at the `[edit protocols isis source-packet-routing srv6]` hierarchy level.

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

- **Support for BGP link-state distribution with SPRING extensions** (ACX Series). Support for BGP link-state extensions export source packet routing in networking (SPRING) and segment-routing topology information to software-defined networking controllers on ACX7100-32C, ACX7100-48L, ACX7509, and ACX7024.

  This feature benefits networks that are moving to SPRING but also have RSVP deployed and continue to use both SPRING and RSVP in their networks.

  [See Link-State Distribution Using BGP Overview, show route table, and show route command.]

- **Support for EVPN VxLAN ACLs and Policers** (ACX7024). Firewall filtering and policing support on EVPN-VXLAN traffic.

  [See Firewall Filters Overview and Understanding EVPN with VXLAN Data Plane Encapsulation.]

- **Support for Flow-Aware Transport for Pseudowires (FAT) label and Entropy label support for EVPN VPWS and flexible cross-connect (FXC)** (ACX7100-32C, ACX7509, and ACX7024 routers). [See FAT Flow Labels Overview .]

- **Support for EVPN VxLAN ACLs and Policers** (ACX7204). Firewall filtering and policing support on EVPN-VXLAN traffic. [See Firewall Filters Overview andUnderstanding EVPN with VXLAN Data Plane Encapsulation.]

- **Support for IPv6 payload-protocol match condition for management interface in ingress and/or egress direction and loopback interface in egress direction** (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5700, QFX5220, and QFX5230-64CD).

  [See Firewall Filter Match Conditions for IPv6 Traffic.]

- **Support for remote authorization on tacplus servers for locally certificate-based authenticated users** (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220). We support remote authorization on TACPLUS servers for locally authenticated users by using certificate-based authentication.

  NOTE: You must configure the password under `authentication-order` when `password-options` is configured. The feature does not work in a local fallback scenario, because the password is not configured under `authentication-order` for a local fallback scenario.

  To configure remote authorization, include the `tacplus-authorization` option under the `password-options` configuration statement at the `[edit system]` hierarchy level.

  [See password-options.]

- **Support for Segment Routing for IPv6 (SRv6) flexible algorithms in traffic engineering database (TED) and BGP Link State (LS)** (ACX7100-32C, ACX7100-48L. ACX7509 and ACX7024). You cannot define flexible algorithms specifically for either SR-MPLS or SRv6. If a node is participating in a flexible algorithm, it would apply to both SR-MPLS and SRv6 nodes.

  [See How to Configure Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering and BGP Link-State Extensions for Source Packet Routing in Networking (SPRING).]

- **Supported transceivers, optical interfaces, and DAC cables** (ACX7024). Select your product in the Hardware Compatibility Tool (HCT) to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

  [See Hardware Compatibility Tool.]

- **VLAN sensor support** (ACX7100-32C, ACX7100-48L, ACX7509, and ACX7024). Junos telemetry interface (JTI) supports data model **openconfig-vlan.yang** version 3.2.1, which includes VLAN sensor support.

  [See Telemetry Sensor Explorer.]

- **Zero-touch provisioning (ZTP) with IPv6 support on management ports** (ACX7024).

  [See Zero Touch Provisioning Overview.]

What's Changed

IN THIS SECTION

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

EVPN

- **Flow-label configuration status for EVPN ELAN services**—The output for the `show evpn instance extensive` command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

- **Commit error if interconnect and local route distinguishers have the same value**—On EVPN data center interconnect (DCI) gateway devices, if you configure an interconnect RD at the **edit routing-instances *name* protocols evpn interconnect** hierarchy, the interconnect RD must be different from the local RD in the routing instance. If you try to configure the same value for the interconnect RD and the local RD in a routing instance, the device enforces this requirement by throwing a commit error. However, with DCI seamless stitching for EVPN Type 5 routes, you don't see the commit error prior to this release. Starting in this release, the device throws the commit error to enforce this condition for DCI stitching with Type 5 routes.

  [See route-distinguisher.]

- **Specify the UDP source port in a ping overlay or traceroute overlay operation** —In Junos OS Evolved releases prior to 22.4R1, you could not configure the udp source port in a ping overlay or traceroute overlay operation. You may now configure this value in an EVPN-VXLAN environment using hash. The configuration option hash will override any other hash-* options that may be used to determine the source port value.

General Routing

- Before this change the output of a `show task replication logical-system all | display xml validate` command reported an error. After the change the output is correctly formatted with a **logical-system** root tag and no validation error occurs.
- Prior to this change the output of a `show task replication | display xml validate` returned an error of the form "ERROR: Duplicate data element <task-protocol-replication-name>. With this change the XML output is properly structured with no validation errors.
- The Ethernet link fault management process (lfmd) runs only when the link-fault-management protocol is configured.
- The connectivity fault management process (cfmd) runs only when the ethernet connectivity-fault-management protocol is configured.
- In the past inet6flow.0 was not allowed to be a primary rib in a rib-group. Starting with Release 22.3 this is now allowed.

Network Management and Monitoring

- **`operator` login class is restricted from viewing NETCONF trace files that are `no-world-readable` (ACX Series, PTX Series, and QFX Series)**—When you configure NETCONF tracing options at the `[edit system services netconf traceoptions]` hierarchy level and you restrict file access to the file owner by setting or omitting the `no-world-readable` statement (the default), users assigned to the `operator` login class do not have permissions to view the trace file.
- **Support for the `junos:cli-feature` YANG extension (ACX Series, PTX Series, and QFX Series)**—The `cli-feature` YANG extension identifies certain CLI properties associated with some command options and configuration statements. The Junos YANG modules that define the configuration or RPCs include the `cli-feature` extension statement, where appropriate, in schemas emitted with extensions. This extension is beneficial when a client consumes YANG data models, but for certain workflows, the client needs to generate CLI-based tools.

  [See Understanding the Junos DDL Extensions YANG Module.]
- **XML tag in the `get-system-yang-packages` RPC reply changed (ACX Series, PTX Series, and QFX Series)**—The `get-system-yang-packages` RPC reply replaces the `xmlproxy-yang-modules` tag with the `proxy-xml-yang-modules` tag in the XML output.
- **Changes to the NETCONF server's `<rpc-error>` element when the `operation="delete"` operation deletes a nonexistent configuration object (ACX Series, PTX Series, and QFX Series)**—We've changed the `<rpc-`

error> response that the NETCONF server returns when the <edit-config> or <load-configuration> operation uses operation="delete" to delete a configuration element that is absent in the target configuration. The error severity is error instead of warning, and the <rpc-error> element includes the <error-tag>data-missing</error-tag> and <error-type>application</error-type> elements.

Platform and Infrastructure

- **The `ping host | display xml` command produces CLI output without errors (ACX Series and QFX Series)** — In Junos OS Evolved release 23.1R1, the `ping host | display xml` command now produces CLI output formatted in XML.

  [See ping.]

System Management

- When subscribing to the resource path /junos/system/linecard/environment, the prefix for the streamed path at the collector side was displaying as /junos/linecard/environment. This issue is resolved in Junos OS Evolved 23.1R1 and the subscription path and the streamed path match to display /junos/system/linecard/environment.
- When disk usage for the **run** directory is above 85%, ZooKeeper logs and snapshots in the **/run/zookeeper/conf/default/version-2** directory will be deleted if there are more than 3 files, leaving only the 3 most recent files.
- **Label-switched interface (LSI) delay during reboot (ACX Series)**—Rebooting ACX Series routers running Junos OS Evolved with a class-of-service routing-instance configuration might encounter errors due to a delay with the label-switched interface (LSI). LSI state information has been added to the output of the show route instance command to assist in the analysis of such errors.

  [See show route instance.]
- **CPU utilization greater than 100% (ACX Series)** —On ACX Series routers running Junos OS Evolved, the `show system processes` command might report CPU utilization spikes greater than 100%. This kind of CPU utilization is normal behavior, and no user action is required. The CPU utilization spikes represent the sum of individual processor threads and not of the entire system CPU capacity.

  [See show system processes.]
- **Mozilla certification authority (CA) certificates removed (ACX Series, PTX Series, and QFX Series)**—To minimize security risks, Junos OS Evolved no longer includes Mozilla's set of root certificates from various CA operators by default. To use Docker container images from a registry that requires TLS

authentication, you must first save the image as a tar archive on a remote device and then import the contents of the archive on the device running Junos OS Evolved.

[See Running Third-Party Applications in Containers.]

Known Limitations

IN THIS SECTION

Learn about limitations in this release for ACX Series routers.

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

General Routing

- syncE to PTP and syncE to 1pps noise transfer tests fail for frequencies on ACX7100-32C
  - 0.00781 HZ
  - 0.01563 HZ
  - 0.03125 HZ
  - 0.06156 HZ
  - 0.12313 HZ

  PR1611911
- On ACX Series routers, 400G DAC flap might be seen after OIR, FPC restart, device reboot, enabling or disabling interface. PR1618488
- ACX7024: With high scale of L3VPN VRF instances system CPU usage might continue to be high.PR1655310

- Due to low CPU on ACX7024, aggregated Ethernet interface takes longer time to come up after deactivating or activating in multi-D scale set up.PR1677809
- When test agents participating to UDP measurement are close by and time is not well synchronised, reporting results is paused due to measured negative one way delay.PR1684358
- Current DHCP server implementation on ACX7100 uses `set access address-assignment pool` *`pool-name`* `family inet network` configuration to allocate all possible IP's from the subnet. One possible workaround is to adjust the netmask according to either increase or decrease the pool size. Full functionality of authd is not available on ACX7100 as it is a small-size platform.PR1689194
- When testing for Class-C performance using channelised 50g interface speed with 100g and 25g as other port, the average cTE marginally exceeds the Class-C mask.PR1695674
- For ACX platforms, protocol mtu (inet, inet6) need to be configured at the family level otherwise family level mtu is not enabled. For example; `set interfaces interface_name unit 0 family inet mtu` *`mtu value`*. PR1696370
- On all Junos OS Evolved platforms, when remote side is applied with 2x100G configuration, and the device is configured with 100g (default speed also), the device's link and the first channel on the remote end comes up.PR1699652
- ACX7024-T-GM: 1PPS, 2Way performance crosses 100 nsec for multi-path scenario in extreme multi-path conditions.PR1707339

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 Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

Authentication and Access Control

- Remote authentication through console is not supported. Local authentication through console and remote authentication through management interface is supported. This is the expected behaviour on Junos OS Evolved going forward. PR1565251

EVPN

- In order to change the flag, routing instances need to be deactivated and activated.PR1685978

General Routing

- With multiple strict priority queues traffic is not scheduled equally between strict priority queues in some specific scenarios. PR1577035
- G.8275.1- G.8273.2 1PPS cTE performance test might be marginally outside class-C for PTP BC on ACX7100-48L, especially for mixed speed port testing with combinations of 10G / 25G channelized ports and 100G ports. On each reboot, the 1PPS cTE measurement might be within the class-C measurement threshold, or might randomly be out of threshold by a few nanoseconds.PR1607381
- syncE to PTP and syncE to 1pps transient response marginally fails. This happens when the servo gets the initial 100ns jump in one measurement window and the next 100ns in the next measurement window adjusting less initially. PR1611848
- ACX7509- Ungraceful removal (OIR) of FPC or an FPC fault might result in PCIE MAJOR alarm **PCI Uncorrected error on dev 0000:00:03.0** which does not get cleared. PR1620197
- If system is fully scaled across features and firewall is also scaled, CPU consumption is more for a small window of around 5 seconds after every 18 seconds or so. Evo-pfemand collects the scaled firewall statistics for the 5 second window and other applications such as **pfe-cli** tries to execute commands that might fail during it. PR1629342
- On ACX7024 ::Pseudo Wire setup and tear down rate might be low. This is due to system CPU limitation.PR1659593
- When TCP Main and TCP remain attached together on physical interface, it is observed that improper scheduler MAP is gets configured on HQoS physical interface while schedule params modification and bindare performed on same commit. This is a sequence issue from CoSD (Routing Engine) which not guaranteed at Packet Forwarding Engine side. And this applicable for all platforms.PR1664785

- In ACX7509, PTP does not work if we move routing network from Layer 2 to Layer 3 protocol on PTP port. PR1669128
- In a ACX7509 system with dual Routing Engine and FEB, when there is a power fault of master FEB, a switchover should happen and backup Routing Engine and FEB takes up primary role. Post switchover, a VMCore file might be generated in the new backup Routing Engine. This failure should not impact the system uptime. Post Vmcore backup Routing Engine reboots and comes back online. PR1671198
- The QSFP28-DD-2X100GBASE-LR4 links go down on multiple FPC restarts or system reboots due to the optics tx loss. PR1685520
- ACX7509: G.8275.1 FPC Dpll status is incorrectly shown in `timingd gencfg ptp centralized` command. PR1685675
- When `set routing-options transport-class auto-create` command is configured, RPD creates or deletes tables dynamically. There is a flaw in the Delete Flow, which does not delete the table from the kernel, and when the next time rpd adds the same table, the operation is stuck with EEXISTS error, as previous delete was never done. Any subsequent commit resolves this issue.PR1696199
- On Junos OS Evolved Platforms, any UI (user interface) set (configuration, script, license) changes done post software addition are lost after the subsequent reboot.PR1699699
- Only one virtual-gateway-v4-mac and one virtual-gateway-v6-mac is supported system wide. The v4 and v6 MAC might be same or different. PR1708967
- ICCP connection establishment between Junos OS and Junos OS Evolved is not supported.PR1710448
- In ACX7509, when the FEB is ungracefully jacked out first and subsequently plugged in, the newly inserted FEB does not come online after online command is issued for the newly inserted FEB. The FEB stays stuck in onlining state.PR1713885

Platform and Infrastructure

- When using RFC2544 benchmarking within L2VPN services in routing Intances, the operational CLI to display the RFC2544 test results returns an error when the routing-instance value is used.PR1696146

Resolved Issues

IN THIS SECTION

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

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

General Routing

- The .include directives are deprecated, and support for them is removed in a future version.PR1647592
- Some VPLS and L2VPN streams are dropped on the aggregated Ethernet interface after a change in MTU configuration on the aggregated Ethernet interface. PR1671451
- The evo-pfemand core file is observed during the device reboot and interface flaps. PR1677679
- ACX7100-32C port down when configuring speed as 10g in Junos OS Evolved 22.2R1. PR1681560
- PTP or TWAMP does not work on PM50 ports (et-0/0/0 to et-0/0/3 on ACX7024 and all ports on ACX7100-48L) if FEC74 is enabled on that port. PR1684770
- Traffic egressing out of layer 3 interface is affected when configuration changes are made. PR1687260
- Traffic failure when you ping from ACX leaf to it's directly attached host.PR1687842
- jdchpd core file seen with dhcp-snooping persistent configuration. PR1688644

- Unable to appropriately failover the incoming traffic with multiple links. PR1688773
- ACX7509 :: timing PTP does not support Routing Engine switchover. PR1690984
- Inconsistent clear ARP behavior leads to traffic loss to host(s). PR1691524
- Junos OS Evolved::JDI_REG::ACX7100: DHCPv6 relay bindings are not as expected after deactivating and activating interfaces. PR1692278
- LACP and LLDP protocol traffic is not transparently forwarded across the CE devices.PR1692402
- [jdhcpd] multiple DHCP release packets sent to the DHCP server in IRB interface with a virtual gateway address (VGA) with EVPN/MPLS ERB model.PR1693625
- RPD core file is seen after the switchover. PR1694773
- System reboot related log message **SYSTEM_REBOOT_EVENT** might not be displayed in `show log messages | match SYSTEM_REBOOT_EVENT`. PR1696668
- On ACX7509, MACsec MKA sessions might not come up after restarting FPC multiple (4-5) times. The only way to recover from this situation is to reboot the chassis.PR1701941
- On Junos OS Evolved platforms, the traffic impact is seen as the `set system process routing enable` or `set system process routing disable` configuration does not work as expected. PR1702734
- Port-mirroring is not working if hierarchical-scheduler is enabled on the analyzer output port.PR1703567
- ACX7509: We see **FEB 0 si5394m_1 PLL Access Failure alarms and Application evo-pfemand fail** on node RE0 after `request system application app hwdre node re0 restart`. PR1708588

EVPN

- In EVPN-MPLS Mutihoming scenario DF election gets stuck in the preference based state.PR1662954
- In the EVPN-MPLS multihoming scenario, MAC-IP route deletion and addition result in traffic drop. PR1691132

Infrastructure

- `show route forwarding-table destination` command takes long time in a scaled system for non-/32 prefixes. PR1685545

Interfaces and Chassis

- LLDP packet drop is seen when the physical interface (IFD) is configured with flexible-vlan-tagging. PR1689391

Routing Protocols

- BGP or OSPF neighbors do not come up in Junos OS Evolved platforms if IPSEC security associations are used to authenticate the peer. PR1674802

Services Applications

- ACX7100: Spikes in jitter that are larger than 5 ms for UDP plugin in PAA for high bandwidth or small packet size. PR1680309
- Impaired accuracy for delay measurements using PAA. PR1697270

User Interface and Configuration

- Test configuration might fail even though the configuration file has valid configurations. PR1671112

Junos OS Evolved Release Notes for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 Devices

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 23.1R1 for PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 Packet Transport Routers. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

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

To view features supported on the PTX platforms, view the Feature Explorer using the following links. To see which features were added in Junos OS Evolved Release 23.1R1, click the Group by Release link. You can collapse and expand the list as needed.

- PTX10001-36MR
- PTX10003
- PTX10004
- PTX10008
- PTX10016

The following sections highlight the key features in this release.

Class of Service

- **Queue-depth monitoring support for virtual output queues (PTX10003)**—Starting in Junos OS Evolved Release 23.1R1, we support the export of buffer utilization data for CoS virtual output queues (VOQs). Using this feature, you can stream telemetry data for peak queue length. Monitoring this data assists in preventing micro-bursts and high buffer utilization for a given queue. To enable VOQ queue-depth monitoring, configure the monitoring-profile statement at the [edit class-of-service] hierarchy level. To export data to a collector, use the following resource paths:

  - /qos/interfaces/interface/input/virtual-output-queues/voq-interface/queues/state/
  - /junos/system/linecard/qmon-sw

  [See VOQ Queue-depth Monitoring.]

- **Shaping overhead adjustment (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, on PTX10000 Series routers you can adjust the number of overhead bytes from –20 to 124. (The default number of shaping overhead bytes added to a packet

at the egress interface is 20 bytes.) This adjustment determines the shaped session packet length at the physical port level.

[See overhead-accounting.]

EVPN

- **Support for EVPN-MPLS (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, you can configure EVPN-MPLS on PTX Series routers. EVPN-MPLS is a solution that extends Layer 2 VPN services over an MPLS network.

  [See EVPN Overview.]

- **Automatically derived ESI support on EVPN-MPLS (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, you can configure multihomed devices in an EVPN-MPLS network to automatically generate the ESI from the:

  - System ID and administrative key on the remote customer edge (CE) device (partner).
  - Locally configured MAC and local discriminator values.

  [See Other Methods to Auto-Derive the ESI.]

- **Determine IRB interface state changes based on local and remote connectivity states in EVPN fabrics (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700)**—Starting in Junos OS Evolved Release 23.1R1, the provider edge (PE) devices in an EVPN fabric consider the following factors when determining the state (up or down) of an L3 integrated routing and bridging (IRB) interface. These factors apply to an L3 IRB interface that is associated with a bridge domain or a VLAN in an EVPN instance (EVI).

  - Associated local L2 interface states

    To customize the L2 interface name and other parameters that the device uses to compute the IRB interface state, configure the `interface-state` statement at the `[edit interfaces irb unit n]` hierarchy level.

  - Remote PE device reachability based on the network isolation state of the bridge domain or the EVI

    The device includes the states of the associated EVPN overlay tunnel interfaces in the network isolation state evaluation.

    To define the parameters that determine when an EVI or a bridge domain is in a network isolation state:

    1. Configure the `network-isolation group group-name` statement at the `[edit protocols]` hierarchy level to define a network isolation profile using the available options.

2. Assign the network isolation group profile to a bridge domain or an EVI using the `network-isolation-profile group` *`network-isolation-group-name`* statement at these hierarchy levels:

   - Bridge domain—`[edit bridge-domain` *`bd-name`* `bridge-options]`
   - EVI—`[edit routing-instance` *`instance-name`* `switch-options]`

[See Determine IRB Interface State Changes from Local and Remote Connectivity States in EVPN Fabrics, interface-state, and network-isolation.]

High Availability

- **MVPN NSR with BGP sharding enabled (PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, we've enabled multicast virtual private network (MVPN) nonstop active routing (NSR) for border gateway function (BGP) sharding.

  [See Understanding BGP RIB sharding and BGP Update IO thread.]

Interfaces

- **400ZR and 400G OpenZR+ Support Enhancements (PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.1R1, we support 400ZR and 400G OpenZR+ optics enhancements on PTX10001-36MR devices. The enhancements include application selection and configuration of target output power. You can view the advertised applications and can switch between the applications.

  See Overview.

Junos Telemetry Interface

- **Number of configurable BMP monitoring stations increases to a maximum of eight (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, and QFX5220)**—Starting in Junos OS Evolved Release 23.1R1, Junos telemery interface (JTI) delivers initial sync and ON_CHANGE BGP routing information base (also known as routing table) statistics by using remote procedure calls (gRPC) or the gRPC network management interface (gNMI) from a device to an outside collector for a maximum of eight BMP monitoring stations.

- **Non-default VRF support for streaming FIB anomalies (PTX10008 and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, Junos telemetry interface (JTI) introduces support for non-default virtual routing and forwarding (VRF) user route Third-Party Attachments (TPAs). TPAs provide more information about an error and are added on top of the errored object. In Junos OS Evolved, the Packet Forwarding Engine (known as a producer) attaches or clears TPAs to the routes during route installation or asynchronously at a later time. The forwarding information base telemetry daemon (FIBTB) (FIB is also known as forwarding table) is an error consumer. The FIBTD listens to and reports

anomalies in a native model to the collector. In this way, the system maintains accurate state information.

The existing FIB anomaly feature supports only default VRF user route TPAs for streaming anomalies. This feature supports non-default VRF user route TPAs.

Use Juniper proprietary remote procedure call (gRPC) services to stream or export ON_CHANGE FIB statistics to an outside SDN collector using the following sensors:

/state/system/anomalies/fib/route-tables/route-table[name]/route/inet/inet-entry[prefix]/status

/state/system/anomalies/fib/route-tables/route-table[name]/route/inet/inet-entry[prefix]/timestamp

/state/system/anomalies/fib/route-tables/route-table[name]/route/inet6/inet6-entry[prefix]/status

/state/system/anomalies/fib/route-tables/route-table[name]/route/inet6/inet6-entry[prefix]/timestamp

[See Error TPAs for Route Installation for information about TPAs. See Telemetry Sensor Explorer for sensor information.]

- **OpenConfig QoS configuration and operational state sensors (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, we introduce support for OpenConfig QoS forwarding classes, classifiers, and rewrites. Support includes configuration and streaming of operational state data.

  [For OpenConfig configuration mappings, see Mapping OpenConfig QoS Commands to Junos Configuration. For state sensors, see Telemetry Sensor Explorer.]

- **Platform sensors for chassis components (ACX7509 and PTX10001-36MR)**—Starting in Junos OS Evolved Release 23.1R1, Junos telemetry interface (JTI) supports operational state sensors for chassis backplane, fabric, fan, power supply, and storage components based on the OpenConfig model **openconfig-platform.yang**. Telemetry data for each component is streamed under the path `/components/component[name=`*`component-name`*`]/`.

  [See Telemetry Sensor Explorer.]

- **Segment routing telemetry for OSPFv2 (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Junos OS Evolved Release 23.1R1 introduces support for the collection and streaming of telemetry data for segment routing (SR) with the OSPFv2 protocol. Statistics are recorded for the Source Packet Routing in Networking (SPRING) traffic per interface, per link aggregation group, and per segment identifier. Support includes OpenConfig and native Junos sensors. To enable collection and export of SR statistics, include the `sensor-based-stats` statement at the `[edit protocol ospf source-packet-routing]` hierarchy level.

[See Telemetry Sensor Explorer for OpenConfig sensors and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface) for native Junos sensors.]

- **VOQ buffer utilization monitoring (PTX10003)**—Starting in Junos OS Evolved Release 23.1R1, we support the export of buffer utilization data for CoS virtual output queues (VOQs). Using this feature, you can stream telemetry data for peak queue length. Monitoring this data assists in preventing micro-bursts and high buffer utilization for a given queue. To enable VOQ queue-depth monitoring, configure the `monitoring-profile` statement at the `[edit class-of-service]` hierarchy level. To export data to a collector, use the following resource paths:

  /junos/system/linecard/qmon-sw/

  /qos/interfaces/interface/input/virtual-output-queues/voq-interface/queues/state/

  [See VOQ Queue-depth Monitoring for information about VOQs. See Telemetry Sensor Explorer for sensor information.]

Licensing

- **Support to trigger license alarm at configured time interval (PTX Series and QFX Series)**—Starting in Junos OS Evolved Release 23.1R1, you can set the time interval at which you want to trigger alarms for features that do not have licenses installed.

  To set the alarm log frequency, use the command `log-frequency` at the `set system license` hierarchy.

  [See Managing Licenses.]

- **MACsec software license feature support (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, soft enforcement of the Media Access Control security (MACsec) license is enabled on listed devices. The license entitles you to:

  - Enable MACsec.
  - Use the available MACsec license bandwidth, which is the sum of the bandwidth associated to each installed MACsec license key.

  In case the feature is enabled without the MACsec license, same will be reported through the alarms, syslog and warning messages on CLI.

  The MACsec bandwidth is tracked at the chassis level, and the configured or default port speed is used as a measure of bandwidth units consumed. An IFD is considered as using the MACsec bandwidth license when:

  - The IFD has MACsec enabled.
  - The IFD admin state is up.
  - At least one user-configured IFL is on the IFD.

[See Flex Software License for PTX Series Routers.]

MPLS

- **OAM support for labeled IS-IS and labeled OSPF flexible algorithm SR paths (PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 23.1R1, Junos OS supports the following Operation, Administration, and Maintenance (OAM) capabilities for labeled IS-IS Flexible Algorithm (flex algo) segment routing paths:

  - IPv4 and IPv6 MPLS ping
  - IPv4 and IPv6 MPLS traceroute
  - Equal-cost multipath (ECMP) traceroute

  Junos OS Evolved also supports IPv4 MPLS ping and IPv4 MPLS traceroute for labeled OSPF flex algo segment routing paths. The OAM functionality is used to detect data plane failures in segment routing paths for the purposes of fault detection and isolation.

  To enable these OAM capabilities, we've introduced the `algorithm` option in the following commands:

  - `ping mpls segment routing isis` *`fec`* `algorithm` *`algorithm-id`*
  - `ping mpls segment routing ospf` *`fec`* `algorithm` *`algorithm-id`*
  - `traceroute mpls segment routing isis` *`fec`* `algorithm` *`algorithm-id`*
  - `traceroute mpls segment routing ospf` *`fec`* `algorithm` *`algorithm-id`*

  [See ping mpls segment routing isis, ping mpls segment routing ospf, traceroute mpls segment-routing ospf, and traceroute mpls segment-routing isis.]

- **Include IGP metric to RSVP routes using conditional metric (MX10004, MX10008, MX10016, PTX1000, PTX10002, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, you can include the IGP metric to RSVP routes that use the conditional metric. Including the IGP metric helps to preserve its value for use in certain use cases, such as in calculating the BGP MED.

  [See Preserving the IGP metric in RSVP LSP routes and include-igp-metricinclude-igp-metric.]

- **Support to report path optimization and computed metrics in PCEP (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, QFX5230-64CD, and QFX5700)**—Starting in Junos OS Evolved Release 23.1R1, we report PCEP path optimization metrics (`IGP`, `TE`, and `delay`) and computed metrics for RSVP and segment routing–traffic engineering (SR-TE) label-switched paths (LSPs).

  To configure the interior gateway protocol (IGP), traffic engineering, and path delay optimization metrics for RSVP LSPs, include the `metric-type` *`igp|te|delay|delay minimum`* CLI statement at the `[edit protocols mpls label-switched-path` *`lsp-name`*`]` hierarchy level.

To configure the optimization metrics for SR-TE LSPs, include the `metric-type` *`igp|te|delay|delay minimum`* CLI statement at the `[edit protocols source-packet-routing compute-profile` *`compute-profile-name`*`]` hierarchy level.

[See Reporting Path Optimization Metrics in PCEP.]

Network Management and Monitoring

- **Minimize system reboot systemd console logging (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220)**—Starting in Junos OS Evolved 23.1R1, we support writing all system application start and application stop messages to the console, not just the error messages. During system reboot, by default, the systemd process logs messages for applications to the journal log. By default, application stop and application start status messages are written to the journal log. Only application stop and application start errors are logged to the console. To change this default, configure the `application-status any` statement at the `[edit system syslog console]` hierarchy level.

  [See application-status.]

- **Sample size support in sFlow (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved 23.1R1, you can configure the sFlow sample size of the raw packet header to be exported as part of the sFlow record to the collector. The configurable range of sample size is from 128 bytes through 512 bytes.

  [See Sample-Size.]

- **Support for additional RPCs for the gNOI `CertificateManagement` service (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, we support the following gRPC Network Operations Interface (gNOI) `CertificateManagement` service RPCs:

  - `CanGenerateCSR()`—Query if the target device can generate a certificate signing request (CSR) with the specified key type, key size, and certificate type.
  - `RevokeCertificates()`—Revoke certificates on the target device.

  [See gNOI Certificate Management (Cert) Service.]

- **Logging support for routing engine shell and line card shell (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, and QFX5130-32CD)**—Starting in Junos OS Evolved Release 23.1R1, you can log commands executed from the shell when you configure `set system syslog shell`.

  [See syslog (System) and start shell.]

Routing Policy and Firewall Filters

- **Support for the IPv6 unicast address-specific BGP extended community attribute (PTX10001-36MR, PTX10004, PTX10008, and QFX5130-32CD)**—Starting in Junos OS Evolved Release 23.1R1, we support the IPv6 unicast address-specific BGP extended community attribute. You can configure the VRF route target with the IPv6 extended community. You can encode each IPv6 unicast address-specific extended community as a 20-octet file.

  To accommodate the IPv6 unicast address-specific extended community, set the IPv6 community configuration under the `[edit policy-options]` hierarchy, and set the following configuration statements in the `[edit policy-options community` *`community-name`* `members]` hierarchy:

  - `ipv6-target:<`*`IPv6 unicast address`*`>:`*`operator-defined local values`*
  - `ipv6-origin:<`*`IPv6 unicast address`*`>:`*`operator-defined local values`*
  - `ipv6-extended:`*`type-and-subtype value`*`:<`*`IPv6 unicast address`*`>:`*`operator-defined local values`*

  [See show route detail, show route advertising-protocol, Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions, Understanding How to Define BGP Communities and Extended Communities, ipv6-extended, ipv6-origin, and ipv6-target.]

Routing Protocols

- **IS-IS Maximum LSP Size (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 23.1R1, you can configure the maximum LSP size for IS-IS in the range 512 through 9192 bytes to support advertisement of a higher number of prefixes.

  [See max-lsp-size.]

- **Support for blocking the redistribution from a specific protocol into IS-IS (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 23.1R1, you can block the export policy from redistributing the routes from other protocols into IS-IS. You can block the redistribution using the `set protocol isis no-external-export` protocol statement at the `[set protocols isis]` hierarchy level.

  [See no-external-export (Protocols IS-IS).]

- **Support for preventing IS-IS from entering overload state on reaching prefix limit (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 23.1R1, you can prevent IS-IS from entering the overload state even after reaching the set prefix limit. You can configure the `set protocols isis dynamic-overload no-overload-on-prefix-export-limit` statement at the `[set protocols isis]` hierarchy level.

  [See no-overload-on-prefix-export-limit (Protocols IS-IS).]

- **Support for auto-recovery from IS-IS overload state (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos Evolved OS Release 23.1R1, you can set up IS-IS to automatically exit an overload state. The command prevents IS-IS from flushing all the LSP fragments on overload so that when the fragment space is available, IS-IS automatically exits from the overload. You can configure the auto-recovery using the `set protocols isis dynamic-overload auto-recovery` statement at the `[set protocols isis]` hierarchy level.

  [See auto-recovery (Protocols IS-IS).]

- **Support for S-BFD over EPE SIDs (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, seamless BFD (S-BFD) running between ingress devices and autonomous system boundary routers (ASBRs) can also track BGP egress peer engineering (EPE) segment identifiers (SIDs). This action helps prevent null-route filtering if a BGP EPE SID goes down.

  [See sbfd.]

- **CLI support for BFD echo and echo-lite modes (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, you can configure BFD `echo` mode and `echo-lite` mode through the Junos OS Evolved CLI. When BFD `echo` mode is active, BFD echo packets are transmitted and looped back from a neighboring device to ensure that a forwarding path is available. BFD `echo` mode requires both the local device and neighbor device to support the full BFD protocol, but BFD `echo-lite` mode can function even if the neighbor device doesn't support BFD.

  You can use the following new CLI configuration statements to configure BFD `echo` mode and `echo-lite` mode:

  - `Echo` mode: `set routing-options static route` *`address`* `bfd-liveness-detection echo minimum-interval` *`interval`*
  - `Echo-lite` mode: `set routing-options static route` *`address`* `bfd-liveness-detection echo-lite minimum-interval` *`interval`*

  [See bfd-liveness-detection.]

- **Support for NDP proxy and DAD proxy for multiple interfaces (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, we support Neighbor Discovery Protocol (NDP) and Duplicate Address Detection (DAD) for interface-restricted mode and interface-unrestricted mode. You use existing CLI commands to enable and disable NDP proxy and DAD proxy for interface-restricted mode and interface-unrestricted mode. Note that you cannot configure `interface-unrestricted` and `interface-restricted` options on the same interface simultaneously.

  [See NDP Proxy and DAD Proxy, dad-proxy, and ndp-proxy.]

- **Support for policy-based ORR (ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220)**—Starting

in Junos OS Evolved Release 23.1R1, the policy-based ORR helps you select the paths to advertise to achieve your traffic-engineering requirements. You can specify a subset of the paths as candidate paths for path selection. The existing path-selection algorithms select the best paths, and you can modify the attributes of the selected paths or reject the selected paths. The policy-based ORR can work alone or with IGP-based ORR and add-path.

To enable this feature in BGP peer groups, configure `export <policy>` at the `[protocols bgp group <name> optimal-route-reflection]` hierarchy level.

[See export (Protocols BGP), optimal-route-reflection, and show bgp group.]

Software-Defined Networking (SDN)

- **Static VXLAN support (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS Evolved Release 23.1R1, you can configure static VXLAN tunnels between VXLAN tunnel endpoints (VTEPs) as follows:

  - In routing instances of type `virtual-switch`
  - Globally (for all VLANs) in a routing instance using the `remote-vtep-list` statement at the `[edit routing-instances` *`name`*`]` hierarchy level.
  - For a specified VLAN in a routing instance using the `static-remote-vtep-list` statement at the `[edit routing-instances` *`name`* `vlans` *`vlan-name`* `vxlan]` hierarchy level. When you specify the static VXLAN remote VTEPs at the VLAN level, you must also specify the same VTEPs at the global level in the routing instance.

  To replicate and flood BUM traffic on the static VXLAN tunnels in addition to forwarding unicast traffic, you must also include the `ingress-node-replication` configuration statement at the `[edit routing-instances` *`name`* `vlans` *`vlan-name`* `vxlan]` hierarchy level. With this statement, the device restricts the BUM traffic flooding domain to only those VTEPs mapped to the specified VLAN.

  [See Static VXLAN and static-remote-vtep-list.]

Source Packet Routing in Networking (SPRING) or Segment Routing

- **Support for unnumbered interfaces for IS-IS with SPRING TI-LFA (PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016)**— Starting in Junos Evolved OS Release 23.1R1, we support IS-IS over unnumbered interfaces on point-to-point links with SPRING topology-independent loop-free alternate (TI-LFA). You can configure unnumbered interfaces to share the same subnet across multiple interfaces to conserve IPv4 addresses. Note that we do not currently support unnumbered interfaces for IPv6.

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

Additional Features

We have extended support for the following features to the platforms shown in parentheses:

- **Inline active flow monitoring IPv6 addressing support for IPFIX and version 9 flow collectors** (PTX10001-36MR, PTX10003, PTX10004, and the PTX10008 and PTX10016 routers with the JNP10K-LC1201 or the JNP10K-LC1202 line card). You can configure either IPv4 or IPv6 collectors for each family within a sampling instance; you cannot specify both for the same family. You can specify up to four collectors for each family within a sampling instance.

  [See Understand Inline Active Flow Monitoring.]

- **MVPN feature support with sharding** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5700).We support the following MVPN features:

  - With sharding, the information required by MVPN, such as inactive L3vpn routes, primary, and secondary details for auto-exported routes were only in the shard thread due to which some MVPN features were only partially supported earlier.

  - As part of the update, these information will be retrieved from shard to main thread. MVPN upstream multicast hop (UMH) mechanisms, sender-based Reverse-path Forwarding (RPF), extranet, and bootstrap (BSR) using auto-export scenarios which depend on these information, will now be in parity/compliant with non-sharding deployments.

  You can use the “show mvpn c-multicast” command to display all the routes on the main thread.

  [See rib-sharding and show mvpn c-multicast.]

- **MVPN support for MLAN, PIM-DIM, and type 7 withdrawal on stream threshold (sharding and non-sharding)** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5700). We support MVPN for MLAN, PIM-DIM, and type 7 withdrawal on stream threshold.

  [See source-active-advertisement, rib-sharding, and show mvpn c-multicast.]

- **Next-Generation MVPN** (PTX10003). We support:

  - Rendezvous-point tree (RPT)-shortest-path tree (SPT) mode.

  - Restart individual PFE instances.

  - Turnaround provider edge (PE) device.

  - RP mechanisms including auto rendezvous point (RP), bootstrap router (BSR), and embedded RP.

  [See Understanding Next-Generation MVPN Control Plane.]

- **Next-generation MVPN Inter-AS option B with or without sharding enabled** (PTX10001-36MR, PTX10004, PTX10008, and PTX10016).

[See BGP-MVPN Inter-AS Option B Overview and Sharding and UpdateIO on cRPD.]

- **Prevent script execution based on current system memory usage** (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700). You can configure the system memory usage threshold above which the device prevents the execution of certain op, event, SNMP, or Juniper Extension Toolkit (JET) scripts. [See Configure Script Start Options.]

- **Real-time Performance Monitoring (RPM) IPv6 source and target address support** (ACX7024, ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700). You can configure IPv6 source and target addresses for RPM. We've also added support for IPv6 addresses to the SNMP RFC2925a MIB control and results tables. For IPv6 RPM probes, you can enable timestamps only in the Routing Engine.

  [See source-address and target.]

- **Support for hybrid mode (Synchronous Ethernet and PTP) over LAG with PTP over Ethernet** (PTX10001-36MR with LC1201-36CD). [See Hybrid Mode.]

- **Support for IPv6 payload-protocol match condition for management interface in ingress and/or egress direction and loopback interface in egress direction** (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5700, QFX5220, and QFX5230-64CD).

  [See Firewall Filter Match Conditions for IPv6 Traffic.]

- **Support for PTP, Synchronous Ethernet, and hybrid mode over LAG** (PTX10001-36MR). We support Precision Time Protocol (PTP), synchronous Ethernet, and hybrid mode over a LAG.

  [See Precision Time Protocol and Hybrid Mode.]

- **Support for QSFP-100G-4WDM40 optics** (PTX10001-36MR).

  [See optics-options.]

- **Support for remote authorization on tacplus servers for locally certificate-based authenticated users** (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220). We support remote authorization on TACPLUS servers for locally authenticated users by using certificate-based authentication.

  NOTE: You must configure the password under `authentication-order` when `password-options` is configured. The feature does not work in a local fallback scenario, because the password is not configured under `authentication-order` for a local fallback scenario.

To configure remote authorization, include the `tacplus-authorization` option under the `password-options` configuration statement at the `[edit system]` hierarchy level.

[See password-options.]

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

What's Changed

IN THIS SECTION

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

EVPN

- **Flow-label configuration status for EVPN ELAN services**—The output for the `show evpn instance extensive` command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport (FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

- **Updated output for show route table**—The output for show route table bgp.evpn.0 now displays L2 service TLV type. Previously, the output displayed the L3 service TLV. PR1694780

- **Commit error if interconnect and local route distinguishers have the same value**—On EVPN data center interconnect (DCI) gateway devices, if you configure an interconnect RD at the `[edit routing-instances name protocols evpn interconnect]` hierarchy, the interconnect RD must be different from the local RD in the routing instance. If you try to configure the same value for the interconnect RD and the local RD in a routing instance, the device enforces this requirement by throwing a commit error. However, with DCI seamless stitching for EVPN Type 5 routes, you don't see the commit error prior to this release. Starting in this release, the device throws the commit error to enforce this condition for DCI stitching with Type 5 routes.

  [See route-distinguisher.]

- **Specify the UDP source port in a ping overlay or traceroute overlay operation** —In Junos OS Evolved releases prior to 22.4R1, you could not configure the UDP source port in a ping overlay or traceroute overlay operation. You may now configure this value in an EVPN-VXLAN environment using hash. The configuration option hash will override any other hash-* options that may be used to determine the source port value.

General Routing

- Before this change the output of a `show task replication logical-system all | display xml validate` command reported an error. After the change the output is correctly formatted with a **logical-system** root tag and no validation error occurs.

- Prior to this change the output of a `show task replication | display xml validate` returned an error of the form "ERROR: Duplicate data element <task-protocol-replication-name>. With this change the XML output is properly structured with no validation errors.

- The Ethernet link fault management process (lfmd) runs only when the link-fault-management protocol is configured.

- Previously, if the system failed to install an interface or hierarchical policer, the PFE crashed due to an assert. Now, the system installs a firewall discard and logs a DFW_HALP_ERR_MSG_POLICER_ADD_FAILED error message. This error message provides the name of the affected policer and the corresponding error code. Relevant policers appear under the interface > unit > family > policer input/output (or) interface > unit > family > input-hierarchical-policer stanzas.PR1701676

- The connectivity fault management process (cfmd) runs only when the ethernet connectivity-fault-management protocol is configured.

- In the past inet6flow.0 was not allowed to be a primary rib in a rib-group. Starting with Release 22.3 this is now allowed.

- **XML tag in the `get-system-yang-packages` RPC reply changed (ACX Series, EX Series, PTX Series, and QFX Series)**—The `get-system-yang-packages` RPC reply replaces the `xmlproxy-yang-modules` tag with the `proxy-xml-yang-modules` tag in the XML output.

- An optics configuration mismatch alarm might be triggered when there is a discrepancy between the configured speed of an interface and the supported speed of the optic. This alarm indicates that the optic installed in the specified FPC is incompatible with the speed configured on the interface.PR1703957

- **Global tunnel termination option disables tunnel termination for all traffic (PTX10000 Series Routers)**—You can use the `set interfaces logical-interface-name unit 'n' family inet/inet6 no-tunnel-termination` command to block VXLAN tunnel termination for the port. Adding the `no-tunnel-termination` option disables tunnel termination for all traffic which the firewall filter would have otherwise allowed you to block termination based on IP addresses.

  [See VXLAN Constraints on PTX10000 Series Routers.]

MPLS

- **Change in display of affinity constraints to hexadecimal values (MX10004, ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10004, PTX10008, and PTX10016)**—Starting in Junos OS release 22.4R1 and Junos Evolved Release 22.4R1, in the output of the `show ted spring-te-policy extensive` operational command, the affinity constraints will be displayed in hexadecimal format instead of decimal.

  [See **show ted spring-te-policy extensive**

Network Management and Monitoring

- **`operator` login class is restricted from viewing NETCONF trace files that are `no-world-readable` (ACX Series, PTX Series, and QFX Series)**—When you configure NETCONF tracing options at the `[edit system services netconf traceoptions]` hierarchy level and you restrict file access to the file owner by setting or omitting the `no-world-readable` statement (the default), users assigned to the `operator` login class do not have permissions to view the trace file.

- **Support for the `junos:cli-feature` YANG extension (ACX Series, PTX Series, and QFX Series)**—The `cli-feature` YANG extension identifies certain CLI properties associated with some command options and

configuration statements. The Junos YANG modules that define the configuration or RPCs include the `cli-feature` extension statement, where appropriate, in schemas emitted with extensions. This extension is beneficial when a client consumes YANG data models, but for certain workflows, the client needs to generate CLI-based tools.

[See Understanding the Junos DDL Extensions YANG Module.]

- **XML tag in the `get-system-yang-packages` RPC reply changed (ACX Series, PTX Series, and QFX Series)**—The `get-system-yang-packages` RPC reply replaces the `xmlproxy-yang-modules` tag with the `proxy-xml-yang-modules` tag in the XML output.

- **Changes to the NETCONF server's `<rpc-error>` element when the `operation="delete"` operation deletes a nonexistent configuration object (ACX Series, PTX Series, and QFX Series)**—We've changed the `<rpc-error>` response that the NETCONF server returns when the `<edit-config>` or `<load-configuration>` operation uses `operation="delete"` to delete a configuration element that is absent in the target configuration. The error severity is error instead of warning, and the `<rpc-error>` element includes the `<error-tag>data-missing</error-tag>` and `<error-type>application</error-type>` elements.

- **Support for export of sFlow samples on the management Ethernet interface (PTX Series)**—You can now export sFlow samples through the management Ethernet interface. Previously, you could only use WAN-facing interfaces to export the samples.

Platform and Infrastructure

- The `ping host | display xml validate` command validates XML without error (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and vMX) — In Junos OS and Junos OS Evolved releases prior to 22.4R2, the `ping host | display xml validate` command results in CRITICAL ERROR: Root tag name mismatch. Expected 'ping-results', got 'run-command'. The command now validates the XML successfully without error.

  [See ping.]

- Prior to this change, devices by default responded only to ARP requests originating from the same subnet. Configure the new CLI option, `respond-out-of-subnet` at the `edit system arp` hierarchy level to allow ARP reply to a request that originates from a different subnet. PR1710699

Routing Protocols

- Prior to this change the output of the `show isis statistics interface <interface_name> | display xml` command used the XML tag "interface-name", which generated an error. With the change the XML output uses the tag "isis-interface-name".PR1712358

System Management

- When subscribing to the resource path /junos/system/linecard/environment, the prefix for the streamed path at the collector side was displaying as /junos/linecard/environment. This issue is resolved in Junos OS 23.1R1 and Junos OS Evolved 23.1R1 and the subscription path and the streamed path match to display /junos/system/linecard/environment.

- When disk usage for the **run** directory is above 85%, ZooKeeper logs and snapshots in the **/run/zookeeper/conf/default/version-2** directory will be deleted if there are more than 3 files, leaving only the 3 most recent files.

- **Mozilla certification authority (CA) certificates removed (ACX Series, PTX Series, and QFX Series)**—To minimize security risks, Junos OS Evolved no longer includes Mozilla's set of root certificates from various CA operators by default. To use docker container images from a registry that requires TLS authentication, you must first save the image as a tar archive on a remote device and then import the contents of the archive on the device running Junos OS Evolved.

  [See Running Third-Party Applications in Containers.]

PKI

- **Deprecating options related to certificate enrollment (Junos)**—Starting in Junos OS Release 23.2R1, we're deprecating earlier CLI options related to Public Key Infrastructure (PKI) to enroll and reenroll local certificate through Simple Certificate Enrolment Protocol (SCEP). The table below shows the Junos CLI commands and configuration statements with the options being deprecated. You can find the same CLI options now available under `scep` option in these commands and statements.

Table 1: Deprecated Junos CLI Options

Junos CLI Commands and Statements	Deprecated Options
`set security pki auto-re-enrollment`	`certificate-id`

`request security pki local-certificate enroll`	`ca-profile`
	`certificate-id`
	`challenge-password`
	`digest`
	`domain-name`
	`email`
	`ip-address`
	`ipv6-address`
	`logical-system`
	`scep-digest-algorithm`
	`scep-encryption-algorithm`
	`subject`

Table 1: Deprecated Junos CLI Options *(Continued)*

Junos CLI Commands and Statements	Deprecated Options

`request security pki node-local local-certificate enroll`	`ca-profile`
	`certificate-id`
	`challenge-password`
	`digest`
	`domain-name`
	`email`
	`ip-address`
	`ipv6-address`
	`logical-system`
	`scep-digest-algorithm`
	`scep-encryption-algorithm`
	`subject`

[See auto-re-enrollment (Security), request security pki local-certificate enroll scep, and request security pki node-local local-certificate enroll.]

User Interface and Configuration

- In Junos OS Evolved releases prior to 22.4R1, the `show system directory-usage` command assumes the current working directory is always **/usr/sbin**. If you want to run the command inside another directory, you must include the full directory path in the command. Starting in Junos OS Evolved Release 22.4R1, this command references the directory you currently have open. The command output displays the absolute path of the directory so you can easily see you are in the correct directory.

  [See system directory-usage.]

Known Limitations

IN THIS SECTION

Learn about limitations in this release for PTX Series routers.

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

MPLS

- Traceroute in MPLS OAM might fail with unreachable in ECMP case when a topology has multiple ECMP paths in each transit router. This is because destination address is not available. Destination address is computed using base address and bitmap index (available for that leg). Junos OS Evolved currently supports 64 bitvector size. Each transit ECMP legs consumes available bitmap indexes in the echo request packet. When all the bitmap indexes are consumed by the previous transit routers and ECMP legs, the bitmap indexes are not available for other ECMP legs. Therefore, multipath information TLV bitmap is zero leading to unreachable issue as no destination address is available. PR1699685

Routing Protocols

- When you do not configure `routing-options transport-class fallback none` statement, you do not configure more than 10 transport-classes. PR1648490

User Interface and Configuration

- On all Junos OS Evolved platforms configured with `persist-group-inheritance`, which is enabled by default from Junos OS Evolved Release 19.4R3 and later, might lead to mustd process crash in highly scaled configuration.PR1638847

Open Issues

IN THIS SECTION

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

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

General Routing

- For the MPC10E card line, the IS-IS and micro-BFD sessions do not come up during baseline. PR1474146
- IPsec tunnels are not deleted on disabling the AMS physical interface. PR1613432
- You can configure **per-interface egress** and **per-sid egress SR sensor statistics** using the following CLI commands: `set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link egress` and `set protocols isis source-packet-routing sensor-based-stats per-sid egress`. As a result, the

(pushed) MPLS label length does not include in the output/Tx octets field that gets exported from the sensor.

PR1646799

- After executing the `clear vpls mac-table` command, not all MAC addresses are learnt for some VPLS instances. PR1664694
- Few protocol sessions remain down causing traffic loss in certain prefixes after quick arpd process disable and enable. The system can be recovered from erroneous state by executing `restart routing gracefully` in CLI. PR1665362
- On all Junos OS Evolved platforms, incorrect sensor base telemetry data are collected when you configure multiple SR-TE tunnels with atleast one uncolored, sharing the same single hop segment list.PR1665943
- UDP telemetry might not work when you subscribe to **/junos/system/linecard/intf-exp/** sensor.PR1666714
- New CLI commands addition to support Routing Engine and chassis power-cycle under `request vmhost` hierarchy. PR1686577
- If you enable MVRP on an MSTP enabled interface, the interface is made part of all the existing instances on the switch. PR1686596
- Junos OS Evolved release has a limitation of 255 characters for resource names. Increasing the limit has implications on the CLI output and same changes are needed to be propagated to lower layers where the resources are served from. PR1695980
- When you configure the command `set routing-options transport-class auto-createis`, rpd creates or deletes tables dynamically. There is a flaw in the delete flow, which does not delete the table from the kernel and when rpd adds the same table next time, the operation is stuck with **EEXISTS** error, as previous delete was never done. .PR1696199
- On all Junos OS Evolved platforms supporting MACsec, traffic drop can be seen when you configure MACsec primary and fallback sessions. As a result, there is a higher transmit-delay time of approximately 6 seconds. This is a timing issue and occurs when switching from primary to fallback or vice-versa. When changing the pre-shared-key's Connectivity Association Key (CAK) value in the CLI on the non-key-server side and at the same time key-server generates a new Secure Association Key (SAK) for pre-shared-key due to expiration of sak-rekey timer, that is, sak-rekey and primary to fallback key-switch both occurs at the same time. .PR1698687
- With the BNG Control Plane User Plane Separation (CUPS) product, you observe an incorrect multicast QoS adjustment with **interface-set queuing** when you execute the `show class-of-service interface-set` command on the user plane. PR1714271

- Whenever a new VLAN is added in between previously configured VLANs, existing context-id which already assigned to existing VLAN context, is assigned for the new VLAN. Therefore, we might see incorrect system-id or bridge-id and this might create an issue.PR1717267

Interfaces and Chassis

- The link-local address is not assigned for the loopback interface after the upgrade or the device reboot on all Junos OS Evolved platforms. The impact depends on how the loopback interface is used in the configuration. It can cause a connectivity issue and traffic impact when it is used for the routing process.PR1695502

Junos XML API and Scripting

- L2TP LAC functionality does not work in this release. PR1642991

MPLS

- When an LSR acts as a Point of Local Repair (PLR) as well as a Merge Point (MP) for an LSP during a double failure scenario, the LSR incorrectly originates one or two **PathErr** messages with routing problem (code=24/2) instead of originating PathErr with notify error (code/subcode=25/3). This does not cause any service impact if the ingress LER might not react adversely to routing problem error (code=24/2).PR1713392

Routing Policy and Firewall Filters

- On all Junos OS Evolved PTX Series platforms, when the source class usage (SCU) or destination class usage (DCU) firewall filter is configured without payload protocol or TCP/UDP protocol, it does not work properly and gives errors whenever the filters are added or re-activated. When the error is encountered, the installation of other filters is also impacted.PR1699138

Routing Protocols

- On all Junos OS Evolved platforms, the rpd can crash when protocol independent multicast (PIM), multicast only fast reroute (MoFRR) configuration is present and some network churn event such as continuous interface cost changes, resulting in a change of active and backup paths for equal cost multi-path (ECMP) occurs. There is service impact because of the rpd crash but the system self-recovers until the next crash. PR1676154
- On all Junos OS Evolved platforms, the commit check fails and blocks LFA itself on one instance if you configure loop-free alternate (LFA) and Remote LFA (RLFA), Per-prefix LFA (PPLFA) in the routing-instance and topology independent LFA (TI-LFA) in the primary instance, along with segment routing and node-link-protection with post-convergence.PR1704521

VPNs

- The rpd crash happens when multicast VPN is configured with separate route-targets scenario. PR1700345

Resolved Issues

IN THIS SECTION

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

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

Class of Service (CoS)

- The host outbound traffic drop is seen on all Junos OS Evolved platforms due to a race condition between class-of-service host-outbound-traffic and forwarding-class configuration update. PR1692542

General Routing

- GNOI rpc KillProcess is not supported. PR1655652
- Routes in RIB and FIB table might go out of synchronization on all Junos OS Evolved platforms and causes a traffic impact. PR1658426
- The license might get out of sync between primary and backup Routing Engine. PR1658869
- The aftmand process crashes upon updating the configurations on the FTI tunnels. PR1663417
- Na-grpcd process core file is seen in telemetry services. PR1665516
- Fragment frames errors are seen on the 400G interface. PR1671065
- The traffic loop is observed when you configure ESI on the physical interface. PR1672631
- Routing Engine reboot, rpd crash, etc. can be seen if the volume of zookeeper log is high on Junos OS Evolved platforms. PR1678880
- GNMI: **/components/component[name=*]/state/oper-status** has duplicate entries for FPC and Routing Engine components. PR1679823
- FPC go to fault state with a major alarm - Power Failure on upgrading. PR1682659
- Junos OS Evolved Release PTX Series devices can reboot in a specific scenario. PR1682898
- Major alarm **Application ztp fail on node Re0** appears about one day later post system zeroize. PR1683964
- The evo-aftmand-bt crashes when the traffic is sampled with egress sFlow. PR1685571

- Type elem not added for INTEGRATED_CIRCUIT. PR1685676
- Logical components which do not have EEPROM, return empty-string for leaves hardware-version, part-no and id. PR1685968
- The fibd process crashes when a large number of interfaces are deleted and added back. PR1685995
- The rpd crash is observed when two separate next-hops in rpd map to the same next-hop-index in the kernel. PR1686211
- The rpd process crash is seen when the BGP SR-TE tunnel is marked for deletion. PR1687287
- queue-counters-queued-bytes-rate for network-class is not within the range with CoS mru 9200 configuration. PR1691957
- CBC-FPGA and RE-FPGA firmware upgrades fail. PR1692186
- CM major errors alarm is not raised for pre_init_pll_programming failure with faulty Packet Forwarding Engine on PTX10004/PTX10008/PTX10016 Junos OS Evolved platforms. PR1693511
- Context deadline exceeded observed on while adding next-hop, IPv4. PR1693567
- The fabspoked-pfe process crashes when a FATAL ERROR occurs in the Packet Forwarding Engine. PR1693697
- CM alarm does not trigger for Packet Forwarding Engine going into fault state. PR1693710
- license-check warning reported on backup Routing Engine by commit or commit check. PR1694935
- License key is not installed after upgrade. PR1696879
- BGP sessions getting flapped. PR1697099
- Junos OS Evolved jkey path changed under protocol IS-IS **IS-IS/levels** is missing PR1698192
- PTX10008 Junos OS Evolved: FTC FPGA minimum supported firmware version mismatch alarm gets generated upon re-seating FTC. PR1698209
- PTX10004/ PTX10008/ PTX10016 Junos OS Evolved: SNMP jnxLEDState mib returns 4(red) value even when BITS LED is unlit/off. PR1698919
- One of the BFD state is down after performing BT restart. PR1699323
- EVO_REG:PTX10008:fibd object-info anomalies observed@net::juniper::addrwatch::AddrWatchNotify. PR1704379
- PTX10004/PTX10008/PTX10016 Junos OS Evolved : When offlined FPC is removed from chassis, `show chassis craft-interface` CLI shows **Fail** for removed FPC. PR1706601

- PTX10008 Junos OS Evolved: When PSM (JNP10K-PWR-AC2) is switched off and "OK LED" is unlit off, but jnxLEDState mib shows green(2). PR1708892
- Observed vmcore while executing MTS (scripr profile: ospf_db_protection_mts_001.robot_BRACKLA.... #bad_area_nosemaphore, #uio_dma_buf_ops_release, #task_work_run). PR1711964
- PTX10004/PTX10008/PTX10016 Junos OS Evolved : SNMP trap jnxFruOnline is raised when JNP10K-PWR-DC2 PSM(single power input) is powered off by button press. PR1713462
- On Junos OS Evolved based PTX Series platforms cannot confirm MAC accounting information. PR1716569

Infrastructure

- The `show route forwarding-table destination` command takes long time in a scaled system for non-/32 prefixes. PR1685545
- The incorrect source address will be used in egress packets. PR1696056
- When a syslog is generated and transported to a log collector over an IPv6 connection, processes like eventd might crash. PR1703823
- SYN-ACK and subsequent TCP session packets generated by Routing Engine will have incorrect DSCP value. PR1703955

Interfaces and Chassis

- LLDP packet drop is seen when the physical interface is configured with `flexible-vlan-tagging`. PR1689391
- The link-local address is not generated for loopback interface. PR1695502

Layer 2 Ethernet Services

- The ethernet switching tables are not in sync between two PE devices. PR1686546

MPLS

- LDP IPv6 session fails to come up in dual transport scenario. PR1683410

Network Management and Monitoring

- snmp-subagent generates a core file at #0 0x00007f54fa7a525d. PR1683517
- Interface physical IP address is assigned to SNMP trap source address after chassis restart although you configure source-address lo0 under trap-options. PR1690850

Routing Policy and Firewall Filters

- Error messages are observed while configuring the firewall filter with family inet6 with next-header and no payload-protocol and committing them. PR1674893
- Unified-ISSU issue is seen on Junos OS Evolved release. PR1685262
- The SCU/DCU firewall filter match does not work as expected. PR1699138

Routing Protocols

- PPMD crashes at ppm_destroy_distrib_proto_stats_group_entry ( ). PR1660299
- Source or destination AS fields shows up as 0 in the flow record. PR1670673
- BGP or OSPF neighbors do not come up in Junos OS Evolved platforms if IPsec associations are used to authenticate the peer. PR1674802
- InboundConvergencePending flag is set after Routing Engine switchover. PR1680360
- Telemetry for peer-as does not work. PR1687369

User Interface and Configuration

- FPC ungracefully restarts when cda-bt process crashes. PR1655441

- In Junos OS Evolved `show | display inheritance` does not work correctly for LSPs with whitespace in the name. PR1693630
- The BFD fails to come up when you configure Routing Engine filter with apply-path. PR1698347

Junos OS Evolved Release Notes for QFX5130-32CD, QFX5220, and QFX5700 Devices

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 23.1R1 for QFX5130-32CD, QFX5220-32CD, QFX5220-128C, and QFX5700 switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

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

EVPN

- **MLDv1, MLDv2, and MLD snooping with OISM and AR in EVPN-VXLAN fabrics (QFX5130-32CD and QFX5700)**—Starting in Junos OS Evolved Release 23.1R1, we support optimized intersubnet multicast (OISM) and assisted replication (AR) with Multicast Listener Discovery (MLD) version 1 (MLDv1), MLD version 2 (MLDv2), and MLD snooping for IPv6 multicast data traffic. You can enable MLD and MLD snooping on these devices in EVPN-VXLAN fabrics when you configure the device as:

  - An OISM server leaf or OISM border leaf device.
  - A standalone AR replicator with OISM.

  [See Optimized Intersubnet Multicast in EVPN Networks.]

- **Determine IRB interface state changes based on local and remote connectivity states in EVPN fabrics (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700)**—Starting in Junos OS Evolved Release 23.1R1, the provider edge (PE) devices in an EVPN fabric consider the following factors when determining the state (up or down) of an L3 integrated routing and bridging (IRB) interface. These factors apply to an L3 IRB interface that is associated with a bridge domain or a VLAN in an EVPN instance (EVI).

  - Associated local L2 interface states

    To customize the L2 interface name and other parameters that the device uses to compute the IRB interface state, configure the `interface-state` statement at the `[edit interfaces irb unit n]` hierarchy level.

  - Remote PE device reachability based on the network isolation state of the bridge domain or the EVI

    The device includes the states of the associated EVPN overlay tunnel interfaces in the network isolation state evaluation.

    To define the parameters that determine when an EVI or a bridge domain is in a network isolation state:

1. Configure the `network-isolation group` *`group-name`* statement at the `[edit protocols]` hierarchy level to define a network isolation profile using the available options.

2. Assign the network isolation group profile to a bridge domain or an EVI using the `network-isolation-profile group` *`network-isolation-group-name`* statement at these hierarchy levels:

   - Bridge domain—`[edit bridge-domain` *`bd-name`* `bridge-options]`

   - EVI—`[edit routing-instance` *`instance-name`* `switch-options]`

[See Determine IRB Interface State Changes from Local and Remote Connectivity States in EVPN Fabrics, interface-state, and network-isolation.]

Juniper Extension Toolkit (JET)

- **Use JET Interfaces Service API to reduce operational time of port bounces (QFX5130-32CD, QFX5220, and QFX5700)**—Starting in Junos OS Evolved Release 23.1R1, you can use the JET Interfaces Service API to perform a port bounce. A port bounce is the act of disabling and re-enabling a physical interface. To disable the port without using the CLI, set the `disable` attribute in the Interfaces Service RPC message to 1 for that port. To remove that configuration and re-enable the port, set `disable` to 0.

  When you set `disable` to 1 in the RPC message, the API disables the port regardless of the CLI or API configuration for that port. When you set `disable` to 0, the API deletes the setting of the `disable` attribute from the API configuration, so the configuration for the port reverts to the previous configuration.

  Details of the Interfaces Service API are in the **jnx_interfaces_service.proto** file in the JET package.

  [See Overview of JET APIs.]

Junos Telemetry Interface

- **Number of configurable BMP monitoring stations increases to a maximum of eight (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5130-48C, and QFX5220)**—Starting in Junos OS Evolved Release 23.1R1, Junos telemery interface (JTI) delivers initial sync and ON_CHANGE BGP routing information base (also known as routing table) statistics by using remote procedure calls (gRPC) or the gRPC network management interface (gNMI) from a device to an outside collector for a maximum of eight BMP monitoring stations.

Licensing

- **Support to trigger license alarm at configured time interval (PTX Series and QFX Series)**—Starting in Junos OS Evolved Release 23.1R1, you can set the time interval at which you want to trigger alarms for features that do not have licenses installed.

To set the alarm log frequency, use the command `log-frequency` at the `set system license` hierarchy.

[See Managing Licenses.]

MPLS

- **OAM support for labeled IS-IS and labeled OSPF flexible algorithm SR paths (PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5220)**—Starting in Junos OS Evolved Release 23.1R1, Junos OS supports the following Operation, Administration, and Maintenance (OAM) capabilities for labeled IS-IS Flexible Algorithm (flex algo) segment routing paths:

  - IPv4 and IPv6 MPLS ping
  - IPv4 and IPv6 MPLS traceroute
  - Equal-cost multipath (ECMP) traceroute

  Junos OS Evolved also supports IPv4 MPLS ping and IPv4 MPLS traceroute for labeled OSPF flex algo segment routing paths. The OAM functionality is used to detect data plane failures in segment routing paths for the purposes of fault detection and isolation.

  To enable these OAM capabilities, we've introduced the `algorithm` option in the following commands:

  - `ping mpls segment routing isis` *`fec`* `algorithm` *`algorithm-id`*
  - `ping mpls segment routing ospf` *`fec`* `algorithm` *`algorithm-id`*
  - `traceroute mpls segment routing isis` *`fec`* `algorithm` *`algorithm-id`*
  - `traceroute mpls segment routing ospf` *`fec`* `algorithm` *`algorithm-id`*

  [See ping mpls segment routing isis, ping mpls segment routing ospf, traceroute mpls segment-routing ospf, and traceroute mpls segment-routing isis.]

- **Support to report path optimization and computed metrics in PCEP (PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, QFX5230-64CD, and QFX5700)**—Starting in Junos OS Evolved Release 23.1R1, we report PCEP path optimization metrics (`IGP`, `TE`, and `delay`) and computed metrics for RSVP and segment routing–traffic engineering (SR-TE) label-switched paths (LSPs).

  To configure the interior gateway protocol (IGP), traffic engineering, and path delay optimization metrics for RSVP LSPs, include the `metric-type` *`igp|te|delay|delay minimum`* CLI statement at the `[edit protocols mpls label-switched-path` *`lsp-name`*`]` hierarchy level.

  To configure the optimization metrics for SR-TE LSPs, include the `metric-type` *`igp|te|delay|delay minimum`* CLI statement at the `[edit protocols source-packet-routing compute-profile` *`compute-profile-name`*`]` hierarchy level.

  [See Reporting Path Optimization Metrics in PCEP.]

Network Management and Monitoring

- **Minimize system reboot systemd console logging (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220)**—Starting in Junos OS Evolved 23.1R1, we support writing all system application start and application stop messages to the console, not just the error messages. During system reboot, by default, the systemd process logs messages for applications to the journal log. By default, application stop and application start status messages are written to the journal log. Only application stop and application start errors are logged to the console. To change this default, configure the `application-status any` statement at the `[edit system syslog console]` hierarchy level.

  [See application-status.]

- **Logging support for routing engine shell and line card shell (ACX7100-32C, ACX7100-48L, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, and QFX5130-32CD)**—Starting in Junos OS Evolved Release 23.1R1, you can log commands executed from the shell when you configure `set system syslog shell`.

  [See syslog (System) and start shell.]

Routing Policy and Firewall Filters

- **Support for the IPv6 unicast address-specific BGP extended community attribute (PTX10001-36MR, PTX10004, PTX10008, and QFX5130-32CD)**—Starting in Junos OS Evolved Release 23.1R1, we support the IPv6 unicast address-specific BGP extended community attribute. You can configure the VRF route target with the IPv6 extended community. You can encode each IPv6 unicast address-specific extended community as a 20-octet file.

  To accommodate the IPv6 unicast address-specific extended community, set the IPv6 community configuration under the `[edit policy-options]` hierarchy, and set the following configuration statements in the `[edit policy-options community` *`community-name`* `members]` hierarchy:

  - `ipv6-target:<`*`IPv6 unicast address`*`>:`*`operator-defined local values`*
  - `ipv6-origin:<`*`IPv6 unicast address`*`>:`*`operator-defined local values`*
  - `ipv6-extended:`*`type-and-subtype value`*`:<`*`IPv6 unicast address`*`>:`*`operator-defined local values`*

  [See show route detail, show route advertising-protocol, Understanding BGP Communities, Extended Communities, and Large Communities as Routing Policy Match Conditions, Understanding How to Define BGP Communities and Extended Communities, ipv6-extended, ipv6-origin, and ipv6-target.]

Additional Features

We have extended support for the following features to the platforms shown in parentheses:

- **MVPN support for MLAN, PIM-DIM, and type 7 withdrawal on stream threshold (sharding and non-sharding)** (PTX10001-36MR, PTX10003, PTX10004, PTX10008, QFX5130-32CD, and QFX5700). We support MVPN for MLAN, PIM-DIM, and type 7 withdrawal on stream threshold.

  [See source-active-advertisement, rib-sharding, and show mvpn c-multicast.]

- **Prevent script execution based on current system memory usage** (ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700). You can configure the system memory usage threshold above which the device prevents the execution of certain op, event, SNMP, or Juniper Extension Toolkit (JET) scripts. [See Configure Script Start Options.]

- **Real-time Performance Monitoring (RPM) IPv6 source and target address support** (ACX7024, ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220, and QFX5700). You can configure IPv6 source and target addresses for RPM. We've also added support for IPv6 addresses to the SNMP RFC2925a MIB control and results tables. For IPv6 RPM probes, you can enable timestamps only in the Routing Engine.

  [See source-address and target.]

- **Support for fan module replacement related alarms and temperature-based system shutdown** (QFX5230-64CD).

  [See How to Configure Flexible Algorithms in IS-IS for Segment Routing Traffic Engineering

  and BGP Link-State Extensions for Source Packet Routing in Networking (SPRING).]

- **Support for IPv6 payload-protocol match condition for management interface in ingress and/or egress direction and loopback interface in egress direction** (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5700, QFX5220, and QFX5230-64CD).

  [See Firewall Filter Match Conditions for IPv6 Traffic.]

- **Support for remote authorization on tacplus servers for locally certificate-based authenticated users** (ACX7100-32C, ACX7100-48L, ACX7509, ACX7024, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5700, and QFX5220). We support remote authorization on TACPLUS servers for locally authenticated users by using certificate-based authentication.

  NOTE: You must configure the password under `authentication-order` when `password-options` is configured. The feature does not work in a local fallback scenario, because the password is not configured under `authentication-order` for a local fallback scenario.

To configure remote authorization, include the `tacplus-authorization` option under the `password-options` configuration statement at the `[edit system]` hierarchy level.

[See password-options.]

What's Changed

IN THIS SECTION

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

EVPN

- **Avoid multicast traffic loss on OISM server leaf and border leaf devices in scaled EVPN-VXLAN fabrics (QFX5130-32CD and QFX5700 switches)**—You can configure QFX5130-32CD and QFX5700 switches as optimized intersubnet multicast (OISM) server leaf or border leaf devices in an EVPN-VXLAN fabric. In scaled fabrics with many VLANs, EVPN instances, and multicast streams, you might see multicast traffic loss on these devices due to the limited size of the multicast snooping route tables in the PFE. To avoid this problem on QFX5130-32CD and QFX5700 switches with OISM in scaled environments, we require that you configure the `conserve-mcast-routes-in-pfe` option at the `[edit multicast-snooping-options oism]` hierarchy on these platforms. This option is available only on QFX5130-32CD and QFX5700 switches. Use this option when you configure these devices as server leaf or border leaf devices with OISM. Do not configure this option when you configure these devices as standalone assisted replication (AR) replicators with OISM.

- **Flow-label configuration status for EVPN ELAN services**—The output for the `show evpn instance extensive` command now displays the flow-label and flow-label-static operational status for a device and not for the routing instances. A device with `flow-label` enabled supports flow-aware transport

(FAT) flow labels and advertises its support to its neighbors. A device with `flow-label-static` enabled supports FAT flow labels but does not advertise its capabilities.

- **Commit error if interconnect and local route distinguishers have the same value**—On EVPN data center interconnect (DCI) gateway devices, if you configure an interconnect RD at the **edit routing-instances *name* protocols evpn interconnect** hierarchy, the interconnect RD must be different from the local RD in the routing instance. If you try to configure the same value for the interconnect RD and the local RD in a routing instance, the device enforces this requirement by throwing a commit error. However, with DCI seamless stitching for EVPN Type 5 routes, you don't see the commit error prior to this release. Starting in this release, the device throws the commit error to enforce this condition for DCI stitching with Type 5 routes.

  [See route-distinguisher.]

- **Specify the UDP source port in a ping overlay or traceroute overlay operation** —In Junos OS Evolved releases prior to 22.4R1, you could not configure the udp source port in a ping overlay or traceroute overlay operation. You may now configure this value in an EVPN-VXLAN environment using hash. The configuration option hash will override any other hash-* options that may be used to determine the source port value.

General Routing

- Before this change the output of a `show task replication logical-system all | display xml validate` command reported an error. After the change the output is correctly formatted with a **logical-system** root tag and no validation error occurs.

- In the past inet6flow.0 was not allowed to be a primary rib in a rib-group. Starting with Release 22.3 this is now allowed.

- Prior to this change the output of a `show task replication | display xml validate` returned an error of the form "ERROR: Duplicate data element <task-protocol-replication-name>. With this change the XML output is properly structured with no validation errors.

- The Ethernet link fault management process (lfmd) runs only when the link-fault-management protocol is configured.

- The connectivity fault management process (cfmd) runs only when the ethernet connectivity-fault-management protocol is configured.

Network Management and Monitoring

- **`operator` login class is restricted from viewing NETCONF trace files that are `no-world-readable` (ACX Series, PTX Series, and QFX Series)**—When you configure NETCONF tracing options at the `[edit system services netconf traceoptions]` hierarchy level and you restrict file access to the file owner by setting or omitting the `no-world-readable` statement (the default), users assigned to the `operator` login class do not have permissions to view the trace file.

- **Support for the `junos:cli-feature` YANG extension (ACX Series, PTX Series, and QFX Series)**—The `cli-feature` YANG extension identifies certain CLI properties associated with some command options and configuration statements. The Junos YANG modules that define the configuration or RPCs include the `cli-feature` extension statement, where appropriate, in schemas emitted with extensions. This extension is beneficial when a client consumes YANG data models, but for certain workflows, the client needs to generate CLI-based tools.

  [See Understanding the Junos DDL Extensions YANG Module.]

- **XML tag in the `get-system-yang-packages` RPC reply changed (ACX Series, PTX Series, and QFX Series)**—The `get-system-yang-packages` RPC reply replaces the `xmlproxy-yang-modules` tag with the `proxy-xml-yang-modules` tag in the XML output.

- **Changes to the NETCONF server's `<rpc-error>` element when the `operation="delete"` operation deletes a nonexistent configuration object (ACX Series, PTX Series, and QFX Series)**—We've changed the `<rpc-error>` response that the NETCONF server returns when the `<edit-config>` or `<load-configuration>` operation uses `operation="delete"` to delete a configuration element that is absent in the target configuration. The error severity is error instead of warning, and the `<rpc-error>` element includes the `<error-tag>data-missing</error-tag>` and `<error-type>application</error-type>` elements.

- **Multicast debug information added in EVPN options to request system information command (QFX Series)**—The output from CLI command `request support information evpn-vxlan` now includes additional information to help debug EVPN multicast issues.

  [See request support information.]

Platform and Infrastructure

- **The `ping host | display xml` command produces CLI output without errors (ACX Series and QFX Series)** — In Junos OS Evolved release 23.1R1, the `ping host | display xml` command now produces CLI output formatted in XML.

  [See ping.]

System Management

- When subscribing to the resource path /junos/system/linecard/environment, the prefix for the streamed path at the collector side was displaying as /junos/linecard/environment. This issue is resolved in Junos OS Evolved 23.1R1 and the subscription path and the streamed path match to display /junos/system/linecard/environment.
- When disk usage for the **run** directory is above 85%, ZooKeeper logs and snapshots in the **/run/zookeeper/conf/default/version-2** directory will be deleted if there are more than 3 files, leaving only the 3 most recent files.
- **Mozilla certification authority (CA) certificates removed (ACX Series, PTX Series, and QFX Series)**—To minimize security risks, Junos OS Evolved no longer includes Mozilla's set of root certificates from various CA operators by default. To use Docker container images from a registry that requires TLS authentication, you must first save the image as a tar archive on a remote device and then import the contents of the archive on the device running Junos OS Evolved.

  [See Running Third-Party Applications in Containers.]

Known Limitations

IN THIS SECTION

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

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

General Routing

- On QFX5700, 400G DAC flap might be seen after OIR, FPC restart, device reboot, enabling or disabling interface. PR1618488

- On QFX5220-32CD, evo-pfemand core file is seen after unified ISSU upgrade through automated script. This happens because vendor sdk fails to warmboot and so evo-pfemand generates core files. In this scenario, evo-pfemand restarts and service resumes. There is impact to traffic until evo-pfemand initialisation completes.

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 Evolved defects, use the Juniper Networks online Junos Problem Report Search application.

General Routing

- QFX5700 - Ungraceful removal (OIR) of FPC or an FPC fault might result in PCIE MAJOR alarm **PCI Uncorrected error on dev 0000:00:03.0** which does not get cleared. PR1620197
- Layer 2 related daemons - lacpd, ifmand, and arpd - when patched using JSU might cause the Junos OS Evolved router to not boot up.PR1676132
- On QX5700, MKA session establishment might be delayed by up to 15 seconds after the configuration commit.PR1705117

Resolved Issues

IN THIS SECTION

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

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

General Routing

- We observed traffic loss due to multicast routes exceeding the scale for OISM feature. PR1671901
- Post ZTP, QFX5220 needs a reboot for routes to be seen in mgmt_junos.inet.0 table. PR1672097
- QFX5220: ISSU is not successful on a router loaded with Junos OS Evolved 22.3- SSH failure after unified ISSU.
- JDI-RCT:IPCLOS:QFX5130-32CD:400G DAC link does not come up. Vendor's shell shows speed as 12.4G instead of 400g. PR1680009
- QFX5700/ QFX5130 : EVPN-VXLAN OSPF and BGP sessions over IRB do not work. PR1688681
- Qfx5220-128C: System reboot related log message **SYSTEM_REBOOT_EVENT** might not be displayed in `show log messages | match SYSTEM_REBOOT_EVENT`. PR1696668
- QFX5700 MacSec: MKA Sessions might not come up after restarting FPC multiple times.PR1702520

High Availability (HA) and Resiliency

- Traffic is impacted if GR-ISSU fails. PR1694669

Routing Protocols

- Traffic loss observed due to multicast routes exceeding the scale for OISM. PR1671901

- Configuration check-out failed when applying **irb with inet and inet6** and **inet6.0 static route.**PR1692484

User Interface and Configuration

- QFX5220: ISSU is not successful on a device loaded with Junos OS Evolved 22.3, SSH failure after ISSU. PR1679476
- Show commands might not work after ISSU upgrade. PR1692409

Upgrade Your Junos OS Evolved Software

Products impacted: ACX7024, ACX7100-32C, ACX7100-48L, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, QFX5130-32CD, QFX5220-32CD, QFX5220-128C, and QFX5700.

Follow these steps to upgrade your Junos OS Evolved software:

1. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage: https://www.juniper.net/support/downloads/
2. In the Find a Product box, enter the Junos OS platform for the software that you want to download.
3. Select Junos OS Evolved from the OS drop-down list.
4. Select the relevant release number from the Version drop-down list.
5. In the **Install Package** section, 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 device or to your internal software distribution site.
10. Install the new package on the device.

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

For more information about software installation and upgrade, see Software Installation and Upgrade Overview (Junos OS Evolved). For more information about EOL releases and to review a list of EOL releases, see https://support.juniper.net/support/eol/software/junosevo/.

Licensing

In 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that Juniper Networks has developed over the past several years.

The major components of the framework are:

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

For information about the list of supported products, see Juniper Flex Program.

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you to explore software feature information to find the right software release and product for your network.

  https://apps.juniper.net/feature-explorer/

- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved.

  https://prsearch.juniper.net/InfoCenter/index?page=prsearch

- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms.

  https://apps.juniper.net/hct/home

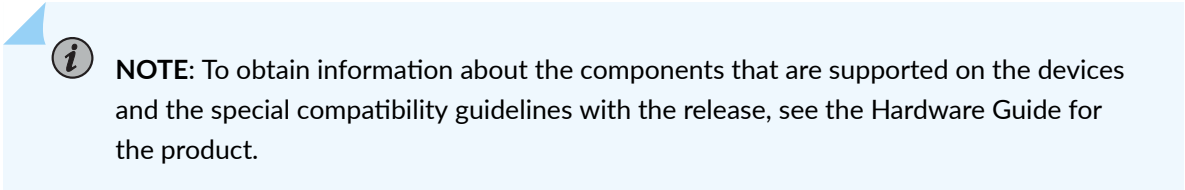

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about Common Criteria, FIPS, Homologation, RoHS2, and USGv6.

  https://pathfinder.juniper.net/compliance/

Requesting Technical Support

IN THIS SECTION

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/

- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit Juniper Support Portal: Case Management, Product Support & More
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see https://support.juniper.net/support/requesting-support/.

Revision History

08 January 2026—Revision 10, Junos OS Evolved Release 23.1R1

11 July 2024—Revision 9, Junos OS Evolved Release 23.1R1

22 March 2023—Revision 8, Junos OS Evolved Release 23.1R1

19 October 2023—Revision 7, Junos OS Evolved Release 23.1R1

10 August 2023—Revision 6, Junos OS Evolved Release 23.1R1

20 July 2023—Revision 5, Junos OS Evolved Release 23.1R1

18 May 2023—Revision 4, Junos OS Evolved Release 23.1R1

20 April 2023—Revision 3, Junos OS Evolved Release 23.1R1

30 March 2023—Revision 2, Junos OS Evolved Release 23.1R1

23 March 2023—Revision 1, Junos OS Evolved Release 23.1R1