Release Notes

Published
2025-09-23

Junos OS Evolved Release 23.4R2

Introduction

Use these release notes to find new and updated features, software limitations, and open issues for Junos OS Evolved Release 23.4R2.

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

Table of Contents

Junos OS Evolved Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 23.4R2 for ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, and ACX7509 devices. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

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

What's Changed

IN THIS SECTION

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

EVPN

- **Updates to syslog EVPN_DUPLICATE_MAC messages**—EVPN_DUPLICATE_MAC messages in the System log (syslog) now contain additional information to help identify the location of a duplicate MAC address in an EVPN network. These messages will include the following in addition to the duplicate MAC address:
  - The peer device, if the duplicate MAC address is from a remote VXLAN tunnel endpoint (VTEP).
  - The VLAN or virtual network identifier (VNI) value.
  - The source interface name for the corresponding local interface or multihoming Ethernet segment identifier (ESI).

  For example: `Feb 27 22:55:13 DEVICE_VTEP1_RE rpd 39839`: EVPN_DUPLICATE_MAC: MAC address move detected for 00:01:02:03:04:03 within instance=evpn-vxlan on VNI=100 from 10.255.1.4 to ge-0/0/1.0.

  For more on supported syslog messages, see System Log Explorer.

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the **edit protocols evpn** hierarchy level. In most use cases, you don?t need to change the default limit. If you want to change the default limit, we recommend that you don?t set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

  [See mac-ip-limit.]

Infrastructure

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

  [See Path MTU Discovery.]

Junos OS API and Scripting

- **<get-trace> RPC support removed (ACX Series, PTX Series, and QFX Series)**—The `show trace application` *app-name* operational command and equivalent `<get-trace>` RPC both emit raw trace data. Because the `<get-trace>` RPC does not emit XML data, we've removed support for the `<get-trace>` RPC for XML clients.

System Management

- **Additional `Upgrade` fields for the `show system applications detail` command (ACX Series, PTX Series, and QFX Series)**—The `show system applications detail` command and corresponding RPC include additional `Upgrade` output fields. The fields provide information about notifications and actions related to various upgrade activities.

  [See show system applications (Junos OS Evolved).]

VPNs

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

  [See min-rate.]

Known Limitations

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

- IPv6 precedence classifier, rewrite CLI is not supported for ACX7000 platforms.PR1587002
- For queues without any guaranteed bandwidth configured, no queue drop counter statistics could be displayed in the case of continuous traffic congestion. This includes those packets getting dropped on queues due to the fact that:
  - Other queue(s) have 100% of bandwidth configured.
  - Other queue(s) are configured as priority queue(s) without any shaper.

  Traffic flow behaviour is as expected per configuration and no impact due to this issue. PR1732194
- While disabling and enabling all the lanes of 400G optics together, carrier transition count on random lanes might get incremented. There is no functional impact.PR1779602

Routing Protocols

- Pruned route entries are not programmed in the case of inclusive tunnels without local receivers. Traffic is dropped in the ingress stage in the PE, hence route resolve does not happen. Thus, this multicast route is not maintained in RPD and the Packet Forwarding Engine. PR1742233

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.

General Routing

- ACX7509 systems have the ability to do automatic fault-based switchovers when traffic impacting faults are seen on the primary Routing Engine or primary FEB. This ability needs to be disabled if the backup Routing Engine is disabled for primary role.

  Suppose a user has disabled RE0 for primary role as follows: `set chassis redundancy routing-engine 0 disabled`This prevents RE0 from becoming the primary on system boot. Instead RE1 becomes the primary. Under this scenario, an automatic fault-based switchover from RE1 to RE0 can still happen. The workaround is to disable automatic fault-based switchovers also as follows: `set chassis redundancy failover disable`. PR1713851.

- We observe core file with rpd with BGP flowspec if secondary-independent-resolution is configured.PR1722715

- On Junos OS Evolved ACX platforms with GRES (Graceful Routing Engine switchover), after performing GRES switchover jdhcpd does not start on the new primary RE (Routing Engine) due to which DHCPv4/v6 (Dynamic Host Configuration Protocol) session binding is lost resulting in traffic loss.PR1740530

- `show pfe statistics traffic` command displays 0 in `Input packets:` and `Output packets:` after Packet Forwarding Engine app (evo-pfemand) restart for some time. Correct values are displayed after the delay.PR1745512

- On Junos OS Evolved ACX platforms configured with EVPN-MPLS (Ethernet VPN-Multiprotocol Label Switching) vlan-based EVPN service or EVPN-VXLAN (Ethernet VPN-Virtual Extensible LAN), device does not respond to ARP/ND (Address Resolution Protocol/Neighbor Discovery) without any vlan-id tag. Due to this, the ARP or ND response packet gets dropped in LAN network and the local host is not able to resolve the ARP.PR1751135

- On ACX7348 and ACX7332 series running Junos OS Evolved, multiple Routing Engine switchover can cause IDEEPROM failure on FPC and PSM, which results in PSM and FPC shown as unsupported.PR1760978

- CFM sessions when scaled on ACX7000 Junos OS Evolved Series can get stuck in OK state on catastrophic events which can result in CCM timeouts at the same time. It can be recovered with CFMD process restartPR1768708

- When the chassis is headless (No Routing Engines), the behaviour of LEDs in that case is not deterministic.PR1769719

- When DHCP `traceoptions` is enabled, there is a possibility that jdhcpd could generate a core file. It is recommended to enable them only for debugging purpose and disable it immediately once debugging is done.PR1771121
- On Junos OS Evolved ACX Series platform in EVPN (Ethernet Virtual Private Network)-VXLAN (Virtual Extensible Local Area Network) scenario in L3VRFs (Virtual Routing and Forwarding) configured with RIB (Routing Information Base) group for route leaks when routing-options multipath is configured the, Type-5 routes advertised through the VRF does not work. There is an impact on the traffic in the VRF when this issue is encountered.PR1773240
- Multihop BFD packets by default uses the network-control queue in Junos OS Evolved ACX Series. This setting remains same despite configuring the host-outbound-traffic to a different forwarding-class.PR1776127
- Junos Evolved based ACX node acting as ASBR+PE (Autonomous System Border Router + Provider Edge) in MPLS L3VPN (Layer 3 VPN) Inter-AS Option B or C does not install VPN label encapsulation in the Packet Forwarding Engine. Traffic is impacted due to this issue.PR1794718
- When multicast packets are transiting ACX7100 devices through VXLAN VTEP or core interface without multicast configurations, errors are seen. Suggested to use DDOS configurations provide with WA. PR1796501
- Under certain conditions when DHCP configuration is completely removed, the session database is left in an unusable state. On subsequent reconfiguration of DHCP, subscriber logins fail. In this case a system reboot is required. PR1816246
- ACX7348::Telemetry :: Few of the sensors do not get streamed for /components/component/. PR1817427
- [evpn-vxlan][dcf12] ACX DC-GW not forwarding traffic received from remote DC-GW after trigger test. PR1817677

Routing Protocols

- BGP passing next-hop path's weight as 0 first and then passing next-hop path's weight as 1 later causes KRT to create 2 different next hops (as weight is now part of next-hop key) and results in installing same route in FIB as primary and backup. PR1745661
- Configuration of a global AS number is necessary when route target filter is enabled. Currently Junos OS does not enforce configuring a global AS number and it has been the behavior for a long time. Many unexpected issues might be seen without a global AS number. It's been a recommended practice to configure a global AS number in the field. PR1783375

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

- G.8275.1- G.8273.2 1PPS cTE performance test might be marginally outside class-C intermittently on ACX7100-48L. PR1607381
- QSFP28-DD-2X100GBASE-LR4 link down on multiple FPC restarts or system reboots. PR1685520
- ACX reports **/psm/0/ hwdre/0/cm/0/ psm_mcu /psm0/psm_cml_cmd_fault** even though the PSM is in working order. PR1700839
- We observe continuous ssh errors on log messages (error: Could not load host key: **/etc/ssh/ssh_host_ec_p521_key**). PR1744354
- Port speed delete is resulting in setting the it to the default speed PR1748138
- We observe spikes in 1pps performance during long run testing. PR1761078
- 100G/400G ports do not come up after performing changes on ACX7332 and ACX7348 platforms. PR1762315
- Any MPLS encapsulated packet with a size less than and equal to 35 bytes gets dropped. PR1766889
- The xinted generates syslog messages `service ssh, accept: Invalid argument (errno = 22)` with high cpu usage. PR1767072

- All Junos OS Evolved platforms - ifmon process 100% utilization. PR1768113
- Unknown unicast IPv4 Traffic received with UDP destination port 8503 is flooded back to Source PE. PR1768729
- IPTV traffic is not forwarded to test agent. PR1771527
- CRC errors observed with SFT-T (740-013111 & 740-027085) Optics. PR1771671
- IPv6 Neighbor Discovery not working resulting in traffic loss. PR1772838
- Traffic forwarding is affected over physical interface if user tries to configure `hierarchical-scheduler` configuration statement on aggregated Ethernet member interface. PR1773980
- [Clocking Solution]:ACX7348: PTP and SyncE does not working properly after Routing Engine switch over. PR1775585
- ACX7348:ACX7332 - Remote interfaces do not go down when chassis is runs headless (with both Routing Engines plugged out from chassis).PR1775785
- In the scaled Layer 2circuit configured with Layer 2 circuit redundancy configuration,traffic drops might be observed. PR1775809
- Aggregated Ethernet load balancing issue on Junos Evolved ACX7000 platforms. PR1775867
- BFD session flaps for member interfaces of the aggregated Ethernet are in multiple line cards. PR1776647
- Traffic drop is seen in EVPN-MPLS aliasing scenario when underlay unilist is modified. PR1776945
- Traffic drops are observed when the aggregated Ethernet link flap in the L2VPN scenario. PR1777608
- clockd core file at **ibdmf::Board::enqueue_event** with restart of FPC multiple times. PR1778163
- Junos OS Evolved: Name resolution does not happen for `show arp` output. PR1778567
- After swtichover with MPLS FRR with VPLS configured, it is observed that traffic to a few VPLS instances is dropped. PR1779466
- Traffic is dropped when `packet-forwarding-options` is configured along with VRF and management VRF configurations. PR1779524
- CLI configuration statement to enable or disable auto recovery of layer 2 learning module in Junos OS Evolved platforms. PR1779797
- Ungraceful FPC removal from chassis causes **hwdre** crash and process stops leading to FPC in fault state. PR1781493

- Multicast traffic is not forwarded to the NG-MVPN core. PR1781735
- Traffic disruption is seen when IRB is present within ERPS protected bridge domain. PR1782190
- ACX7509 shows PLL alarms after an Routing Engine switch. PR1782380
- ACX7024 does not boot after `request system zeroize`.PR1783542
- PTP remains in ACQUIRING state. PR1783545
- Junos OS Evolved: ACX7000 Series: Protocol specific DDoS configuration affects other protocols (CVE-2024-39531). PR1784343
- Policy-based routing does not work as intended when it is configured with `next-ip` or `next-ip6` on all ACX Junos OS Evolved platforms except ACX7024. PR1784909
- STP bridge domain or ERPS protected domain configured on the IRB interface causes traffic to be dropped. PR1784990
- Route leaks between VRFs through the RIB group does not work as expected. PR1786295
- CFM configuration with multiple MEP under same maintenance-association causes evo-pfemand process crash and CFM session might not come up. PR1786395
- Autoneg configuration for 1G optical does not turn on upon inspecting PHY registers. PR1787154
- On ACX7348, after system reboot or Routing Engine switchover, syslog messages with **"ddsType = "RouteCcc" tpName = "BrcmRtCcc""** are seen. These messages don't cause any service or traffic impact. PR1787689
- Interfaces with QSFP-100GBASE-LR4 optics might not come up after software upgrade or system reboot. PR1788848
- The evo-pfemand process crashes on Junos OS Evolved ACX platforms. PR1791199
- Junos ping inet command does not force IPv4 in Junos OS Evolved platforms. PR1792415
- Error messages are populate **getQosRewriteHwMapIdFromIflIndex** and interface ifd queue statistics traffic goes to the wrong queue. PR1793256
- FPCs are not recovered when system is rebooted right after Routing Engine switchover due to evo-pfemand crash. PR1797593
- Traffic drops observed in the Layer 2 circuit or Layer 2 VPN scenario. PR1797839
- Traffic drops are seen on all Junos OS Evolved platforms. PR1798446
- Multiple create or delete ofphysical interface, L3 and MACsec interfaces blocks MACsec traffic on Junos Evolved platforms. PR1800139

- On ACX platforms running Junos OS Evolved, the device gets powered down due to incorrect reading of a temperature sensor, impacting all the services running on the box. PR1801225
- [Junos OS Evolved] **Host 0 Disk 1 Labelled incorrectly** alarm is sometimes set and cleared in 5 seconds PR1801436
- In platform ACX7509 after switchover, STP states re-converge even with NSB enabled. PR1801786
- The rpd process crashes when Routing Instance type changes from L2 to L3 or vice versa. PR1802000
- CB goes into fault state after system Routing Engine power off or on using button press. PR1802508
- On Junos Evolved ACX7000 platforms configured with MPLS tunnel and Storm Control, the Layer 2 or MPLS tunnel-terminated known unicast traffic arriving at the ingress interface assigned with a high drop precedence by the interface-level classifier gets dropped on the interface where storm control profile is active. Due to this, MPLS tunnel traffic might get affected. PR1802525
- [ACX7000] With IPSEC-Ah configuration enabled, OSPF or OSPF3 session does not work. PR1803437
- Some timing features do not work as intended with PLL input and lock failure alarms on Junos OS Evolved platforms after Routing and Control Board jack in and switchover. PR1803481
- ACX7000: Post instance renaming, VPLS MACs stopped exchanging over MPLS core or LSI interface. PR1805586
- ACX7024 Timing : T-GM performance shows 1second time offset. PR1808134
- L2ald-agent generates core file and IRB logical interface stays Hardware-down after deletion of IRB (with virtual-gateway-address configuration) , reading same virtual-gateway-address as IRB address and move back to IRB with same virtual-gateway-address. PR1808779
- OSPFv3 neighborship is not getting established on IRB when mld snooping is enabled for the BD on which IRB is hosted. PR1816540

Infrastructure

- DNS resolution does not work for default instance when name server is reachable through mgmt_junos VRF. PR1766212

Interfaces and Chassis

- On the ACX7024 platform, the tear-down rate is low. This is due to system CPU limitations. PR1659593
- Multiple processes crash when more than 150 VLAN entries are configured in vlan-id-list under aggregated Ethernet logical interface. PR1774222

Layer 2 Features

- MAC learning has been rejected due to mismatch between l2ald and PFE STP index value in VPLS.PR1766991

User Interface and Configuration

- Configuration archival through FTP, does not work with presence of firewall filter on the loopback interface even when FTP ports are allowed. PR1798464

Junos OS Evolved Release Notes for PTX Series

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 23.4R2 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 23.4R2-S1

IN THIS SECTION

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

Hardware

- **New PTX10002-36QDD router (PTX Series)**—In Junos OS Evolved Release 23.4R2-S1, we introduce the PTX10002-36QDD router. It is a fixed-configuration router that features 36 high-density and cost-efficient 800-Gigabit Ethernet ports network ports in a 2-U form factor. With 28.8 terabits per second (Tbps) of throughput, the PTX10002-36QDD is optimally designed for peering, core routing, and infrastructure edge routing roles in cloud provider, service provider, and content provider networks.

  The router supports 2200-W or 3000-W high-voltage HVAC/HVDC and DC power supply units and front-to-back airflow.

  You can channelize the ports on the PTX10002-36QDD and increase the number of interfaces.

  To install the PTX10002-36QDD router and perform initial configuration, routine maintenance, and troubleshooting, see the PTX10002-36QDD Hardware Guide. See Feature Explorer for the complete list of features for any platform.

Table 1: Features Supported on the PTX10002-36QDD

Feature	Description
Chassis	• Support for the following chassis management functionalities: • The presence of two ASIC packages enables you to take a Flexible PIC Concentrator (FPC) online or offline or to restart the FPC without impacting the power to the FPC. • When you connect 3000-watt (W) power supply units (PSUs), the system operates in normal power mode. You can change the operating power mode from normal to power-optimized by using the command `set chassis mode power-optimized`. • The `show chassis fpc` command displays both `PFE` and `PFE-Instance` details. • On the router, when you run the `request chassis fpc` command, you must use `pfe`instead of `pfe-instance` to control the FPC operations. Also, when you run the `request chassis fpc` command, you must commit the command for both the PFEs that are present. [See **Power Mode Management on PTX10002-36QDD**, chassis, request chassis fpc, and show chassis fpc.] • Supports resiliency features to manage fabric faults, including but not limited to: • Auto-heal functionality to recover the faulty link by fixing the errors automatically. • Disables all Packet Forwarding Engines when the number of fabric link errors exceeds four in the system. You can use the existing CLI commands for the fabric management. The following commands display new or different fields in their outputs: • `show chassis fabric fpcs` displays peer `FPC/PFE` details because the Packet Forwarding Engines are directly connected to each other.

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• `show chassis fabric topology` displays only physical link connectivity. [See Chassis-Level User Guide, show chassis fabric fpcs, and show chassis fabric topology.] • Optics EM policy support. We've extended the Junos Environment Monitoring (EM) policy to include optics temperature sensors for PTX10002-36QDD routers. It includes the following features: • The Optics EM policy incorporates periodically polled temperature readings of optical modules in the system to automatically manage the fan speed. • Junos OS Evolved will automatically trigger optics shutdown for 100GbE, 400GbE, and 800GbE optics when the Fire Shutdown threshold is breached. Auto-recovery is not supported for optics disabled by the EM policy. To re-enable the optics, use the `request interface optics-reset` command or perform soft optics insertion and removal (OIR). • EM policy is enabled by default on all 100GbE, 400GbE, and 800GbE optics that are Multi-source Agreements (MSA) compliant and support diag EEPROM with temperature monitoring. This policy is not applicable for loopback optics and direct attach copper (DAC) cables. To disable EM policy or view temperature threshold values, use the following CLI commands: • `set chassis fpc` *`fpc_slot`* `pic` *`pic_slot`* `port port_no no-temperature-monitoring` explicitly disables the EM policy on specific WAN ports. • `show chassis temperature-thresholds` displays the optics temperature threshold values. • `show chassis environment` displays the optics temperature. [See chassis-adc-temperature-sensor.]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Routing Engine support. The fixed-configuration PTX10002-36QDD router supports an inbuilt Routing Engine represented by the model number RE-JNP10002-36QDD in the CLI. The router does not support: the following operational commands: • A pluggable Routing Engine • Graceful Routing Engine switchover (GRES), as the router does not have a redundant Routing Engine • The following operational commands: • `request chassis routing-engine master acquire` • `request chassis routing-engine master release` [See show chassis hardware.] • Routing Engine resiliency. We've enabled Routing Engine resiliency for the faults related to CPU memory and DIMM. The Routing Engine supports fault-handling actions such as logging errors, raising alarms, sending SNMP traps, and providing indication about anerror through the LEDs. [See show system errors active.] • Support for fabric platform resiliency includes resiliency functionality to manage hardware components such as the FPC, PSUs, and fan. [See show chassis power detail, show chassis fpc, and show chassis fan.] • Packet Forwarding Engine resiliency. The software detects, reports, and takes action on Packet Forwarding Engine faults. Actions are taken based on the default configuration or user configuration available for the errors. [See show system errors active.]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Class of service	• Support for class-of-service (CoS) features, including classifiers (behavior aggregate (BA), fixed, and multifield (MF)), rewrite rules, forwarding classes, loss priorities, transmission scheduling, rate control, and drop profiles. [See CoS Features and Limitations on PTX Series Routers.] • Support for priority-based flow control (PFC) watchdog, which detects and mitigates PFC pause storms received for PFC-enabled queues. A table called `jnxCosWatchdogTxQueueTable` is added to the SNMP class-of-service (CoS) MIB to show statistics for transmitting PFC queues related to the PFC watchdog. Table entries are indicated by `jnxCosWatchdogTxQueueEntry` and contain the following objects: • `jnxCosWatchdogIfIndex`—The index of an interface on which PFC and PFC watchdog are enabled. • `jnxCosWatchdogTxQueueId`—The ID of the queue of the PFC-enabled interface. • `jnxCosWatchdogTxQueueRecoveredCount`—The number of times a queue recovered after a PFC pause storm. • `jnxCosWatchdogTotalPktDrop`— The total number of packets dropped due to PFC pause storm mitigation since the device was started. • `jnxCosWatchdogLastPktDrop`—The number of packets dropped due to the last PFC pause storm. [See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved and PFC Watchdog.] • Support for importing existing classifier and rewrite rules to form new rules. • Support for priority-based flow control (PFC) at Layer 3 for untagged traffic and explicit congestion notification (ECN).

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	[See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic and CoS Explicit Congestion Notification.] • Queue-depth monitoring support for virtual output queues. Virtual output queue (VOQ) queue-depth monitoring, or latency monitoring, measures peak queue occupancy of a VOQ. Junos OS Evolved supports VOQ queue-depth monitoring to report peak queue length for a given physical interface for each Packet Forwarding Engine. [See VOQ Queue-depth Monitoring.] • Support for export of physical interface queue statistics to an outside collector. Use UDP (native) streaming, remote procedure call (gRPC) services, or gRPC network management interface (gNMI) services by using the sensor **/junos/system/linecard/interface/queue/**. Each physical interface has eight queues. The following counters are exported as part of this sensor for all configured physical interfaces: • Transmitted packets and transmitted bytes • Red drop packets and bytes • Tail drop packets and bytes This feature includes zero suppression support. It does not include support for summed up counters on aggregated Ethernet (ae) interfaces. [See sensor (Junos Telemetry Interface) and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).] • Hierarchical CoS support. The router supports up to four levels of scheduling on an interface (physical interfaces, logical interface sets, logical interfaces, and queues). The router does not support hierarchical CoS on integrated routing and bridging (IRB) or aggregated Ethernet interfaces. Also, hierarchical CoS schedulers should not include buffer or drop profile configurations.

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	To enable hierarchical scheduling, set `hierarchical-scheduler` at the [`edit interfaces` *`interface-name`*] hierarchy level. [See Hierarchical Class of Service in ACX Series Routers.] • Support for classification override configured under a forwarding policy. [See CoS Features and Limitations on PTX Series Routers and Overriding the Input Classification.]
Dynamic Host Configuration Protocol	• DHCPv4 Relay Agent and DHCPv6 Relay Agent are supported. Features included are: • DHCP Relay: Layer 3 (L3) interfaces • DHCP Relay: Option 82 for Layer 2 VLANs • DHCP Relay: Option 82 for L3 interfaces • Extended DHCP Relay Agent • Virtual router-aware DHCP (VR-aware DHCP) [See Extended DHCP Relay Agent Overview.]
Hardware	• Supported transceivers, optical interfaces, and DAC cables. Select your product in the Hardware Compatibility Tool to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available. [See Hardware Compatibility Tool .]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
High availability and resiliency	• BFD support, including: • Distributed BFD and BFD-triggered local repair (BFD authentication is not supported.) • Independent micro BFD sessions enabled on a per-member link basis for a LAG bundle • Inline BFD [See Understanding BFD .] • Support for IP-over-IP encapsulation to facilitate IP overlay construction over an IP transport network.—An IP network contains edge devices and core devices. To achieve higher scale and reliability among these devices, use an overlay encapsulation to logically isolate the core network from the external network that the edge devices interact with. Static configuration or a BGP protocol configuration is used to distribute routes and signal dynamic tunnels. The dynamic-tunnels configuration creates IP-over-IP encapsulation-only tunnels in the Packet Forwarding Engine. The following are not supported: • Dynamic tunnel de-encapsulation operation • Next-hop-based statistics for dynamic tunnels • IP fragmentation at tunnel start point and path MTU discovery for IPv4/IPv6 [See Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation .] • Support for VRRP. The following features are not supported for VRRP on Junos OS Evolved: • ISSU

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Proxy ARP • MC-LAG • Distribution support on aggregated Ethernet interfaces • IRB • Inline delegation [See Understanding VRRP .]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Interfaces	• Interface support. The PTX10002-36QDD supports multiple port speeds and various channels under each port. The router supports a maximum port speed of 400 Gbps in low power mode. In standard power mode, it supports a port speed of 800 Gbps, If a port speed is configured (or a port speed is determined by default) without configuring `number-of-sub-ports` (at [`edit interfaces` *`interface-name`*] hierarchy level), the port operates in non-channelized mode. [See Port Speed on PTX Routers.] • 400G-ZR and 400G-ZR+ support enhancements. We support 400G-ZR and 400G-ZR+ optics enhancements on the PTX10002-36QDD. The enhancements include application selection and configuration of target output power. You can view the advertised applications and switch between the applications. [See Features of 400ZR and 400G OpenZR+.] • Support for performance monitoring and TCA. We support performance monitoring for the PTX10002-36QDD optical transceiver modules. The current and historical performance monitoring metrics are accumulated into 15-minute and 1-day interval bins. You can view the metrics using the `show interfaces transport pm` command and manage optical transport links efficiently. [See show interfaces transport pm.] • Support for timing and synchronization. The PTX10002-36QDD supports Synchronous Ethernet compliant with the following ITU recommendations: • G.8262/G.8262.1—Specifies timing characteristics of Synchronous Ethernet equipment clock (EEC). • G.8264—Describes the Ethernet Synchronization Message Channel (ESMC). [See Synchronous Ethernet Overview.]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Support for load balancing under the `[edit forwarding-options enhanced-hash-key]` hierarchy. Load balancing includes: • GRE key inclusion for transit IPv4 and IPv6 traffic • IP Layer 3 fields • IP Layer 4 fields • IPv6 flow label inclusion • MPLS labels • MPLS port data • MPLS pseudowire traffic • Tunnel endpoint identifier (TEID) inclusion in GPRS tunneling protocol (GTP) packets • RSVP-TE load balancing in proportion to LSP bandwidth [See enhanced-hash-key.] • Support for 128-way equal-cost multipath (ECMP) routing for MPLS transit cases. The following features do not support 128-way ECMP: • Multicast • P2MP • MC-LAG • Weighted unilist • Consistent hashing • Link protection (MPLS)

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Adaptive load balancing • Class-based forwarding • Support for 256-way ECMP. You can configure a maximum of 256 equal-cost multipath (ECMP) next hops for external BGP (EBGP) peers. This feature increases the number of direct BGP peer connections, which improves latency and optimizes data flow. However, we support 128 ECMP next hops for MPLS routes. Note that we do not support consistent load balancing (consistent hashing) for IPv4 or IPv6 with this feature. [See Understanding BGP Multipath.] • Support for FTI-based encapsulation and de-encapsulation of IPv4 and IPv6 packets. You can configure IP-IP encapsulation and de-encapsulation on flexible tunnel interfaces (FTIs). The default mode is loopback encap mode. Use the `bypass-loopback` statement at the `[edit interfaces fti` *`number`* `unit` *`logical-unit-number`* `tunnel encapsulation ipip]` hierarchy level to change into flattened encap mode to achieve line-rate performance. [See Tunnel and Encryption Services Interfaces User Guide for Routing Devices.] • Support for configuring UDP tunnel encapsulation on FTIs. You can configure encapsulation by using the `tunnel encapsulation udp source` *`address`* `destination` *`address`* statements at the `[edit interfaces fti unit` *`unit`*`]` hierarchy level. Keep in mind the following when configuring this feature: • Adding `tunnel-termination` makes the tunnel a de-encapsulation-only tunnel and encapsulation is disabled. • Specifying both the source and destination address is mandatory when you do not configure `tunnel-termination`. • Configuring a variable prefix mask on the source address is not allowed.

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	[See encapsulation (interfaces-fti).] • **GRE tunnel encapsulation using loopback-based interface.** You can configure GRE tunnel encapsulation on flexible tunnel interfaces (FTIs) using the loopback interface. Configure encapsulation by using the `tunnel encapsulation gre source` *`address`* `destination` *`address`* statements at the `[edit interfaces fti0 unit` *`unit`* `]` hierarchy level. [See encapsulation (interfaces-fti).] • Support for GRE tunnel de-encapsulation using FTIs. Flexible tunnel interfaces (FTIs) support GRE tunnel de-encapsulation. When you enable the `tunnel-termination` statement at the `[edit interfaces fti0 unit` *`unit-number`*`]` hierarchy level, tunnels are terminated on the WAN interface before any other actions—such as sampling, port mirroring, or filtering—are applied. [See Tunnel and Encryption Services Interfaces User Guide for Routing Devices.] • Support for configuring MPLS protocols over FTI tunnels, thereby transporting MPLS packets over IP networks that do not support MPLS. Generic routing encapsulation (GRE) and UDP tunnels support the MPLS protocol for both IPv4 and IPv6 traffic. You can configure encapsulation and de-encapsulation for the GRE and UDP tunnels. To allow the MPLS traffic on the UDP tunnels, include the `mpls port-number` statement at the `[edit forwarding-options tunnels udp port-profile` *`profile-name`*`]` hierarchy level. To allow the MPLS traffic on the GRE tunnels, include the mpls statement at the `[edit interfaces fti0 unit` *`unit`* `family]` hierarchy. [See Flexible Tunnel Interfaces Overview.]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Junos telemetry interface	• JTI support for Packet Forwarding Engine sensors for usage, network processing unit (NPU) memory, NPU utilization, and pipeline NPU and ASIC. Using the Junostelemetry interface (JTI), you can export statistics using remote procedure call (gRPC) services, gRPC Network Management Interface (gNMI) services, and UDP transport. Use these sensors: • `/junos/system/linecard/packet/usage/` • `/junos/system/linecard/npu/memory/` • `/junos/system/linecard/npu/utilization/` • `/components/component/integrated-circuit/state/` • `/components/component/integrated-circuit/pipeline-counters/` For pipeline sensors, the four packet and drop counter categories are interface, lookup, queuing, and host interface. [See Junos YANG Data Model Explorer.] • JTI support for platform sensors. Using Junos telemetry interface (JTI), you can export platform-specific software and chassis component statistics using remote procedure call (gRPC) services, gRPC Network Management Interface (gNMI) services, and UDP transport. Use these sensors: • `/junos/system/cmerror/` • `/junos/system/linecard/` • `/components/components/` • `/system/alarms/` • `/state/interfaces/`

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• `/state/chassis/` [See Junos YANG Data Model Explorer.]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Layer 2 features	• Support for flow-aware transport (FAT) for pseudowires labels on ingress routers, with parsing that includes all the payload fields in the hash calculation. These flow labels are supported: • L2circuit, LDP-signaled pseudowires • L2VPN, BGP-signaled pseudowires • L2VPN with FEC129 (BGP autodiscovery) [See **flow-label-receive** and **flow-label-transmit**.] • Support for VLAN tag manipulation: pop, push, and swap. [See **Configuring an MPLS-Based VLAN CCC with Pop, Push, and Swap and Control Passthrough**.] • Support for virtual circuit connection verification (VCCV) protocol, which transfers control packets from one provider edge (PE) router to another PE router by creating a separate channel in the pseudowires. The pseudowires set up signaling peers and use a control word to maintain proper sequencing of pseudowire packets over the packet-switched network. [See **BFD Support for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS**, **Configuring BFD for VCCV for Layer 2 Circuits**, **MPLS Pseudowires Configurationshow ldp database**, and **show route instance**.] • Support for inner VLAN transparency. We support the pop, push, swap, pop-pop, pop-swap, swap-push,push-push, and swap-swap operations on port-based and VLAN-based Metro Ethernet Forum (MEF) Layer 2 services. VLAN transparency refers to preserving inner VLANs in the packet that are not subject to manipulation and are not used for forwarding. Based on the scenarios, VLAN transparency works on up to four VLAN tags. [See **Understanding VLAN Manipulation (Normalization and VLAN Mapping) on Ethernet Services**.] • Support for the following protocols:

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	1. LAG (aggregated Ethernet) 2. LACP 3. LLDP

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Layer 3 features	• Support for the following Layer 3 forwarding features on the router: • IPv4 • IPv6 • MPLS • LAG • ECMP • MTU checks • ICMP • OSPF • IS-IS • ARP • NDP • BGP • BFD • LACP • LDP • RSVP • LLDP • VRF-lite • TTL expiry • IP options

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• IP fragmentation • DDoS

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
MACsec	• MACsec support in static CAK mode on physical interfaces with dynamic power management.—Media Access Control Security (MACsec) is an industry-standard security technology that provides secure communication for traffic on Ethernet links. This device supports MACsec in static connectivity association key (CAK) mode.This device supports MACsec on physical interfaces to enable you to secure your network using any of the following encryption types: • GCM-AES-128 • GCM-AES-256 • GCM-AES-XPN-128 • GCM-AES-XPN-256 This device supports the following MACsec features: • Configurable security association key (SAK) rekey period • MACsec Key Agreement (MKA) protocol fail open mode • Preshared key (PSK) chains and hitless rollover • PSK password encryption using master password • Fallback PSK • Extended packet numbering (XPN) • Jumbo frames [See Understanding Media Access Control Security (MACsec).] • MACsec dynamic power management support. Use MACsec to secure your network with the knowledge that your device is working to optimize power usage. To save power, the device dynamically powers MACsec blocks on and off based on the MACsec configuration.There might be minimal traffic loss during the power block transition. [See Understanding Media Access Control Security (MACsec).]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
MPLS	• Support for MPLS FRR. MPLS fast reroute (FRR) provides faster convergence time (less than 50 milliseconds) for RSVP tunnels. The Routing Engine creates backup paths and the Packet Forwarding Engine installs the backup-path labels and next hops. [See Fast Reroute Overview.] • Support for MPLS features,including: • CLI support for monitoring MPLS label usage • Inline MPLS and IPv6 lookup for explicit null • 32,000 transit LSPs • Explicit null support for MPLS LSPs • MPLS Label Block Configuration • MPLS over untagged Layer 3 interfaces • MPLS OAM - LSP ping • JTI: OCST: MPLS operational state streaming (v2.2.0) • 2000 ingress LSP support • 2000 egress LSP support • Entropy label support • MPLS: JTI: Junos telemetry interface MPLS self-ping, TE++, and misc augmentation • LDP, including: • Configurable label withdraw delay • Egress policy

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Explicit null • Graceful restart signaling • IGP synchronization • Ingress policy • IPv6 for LDP transport session • Strict targeted hellos • Track IGP metric • Tunneling (LDP over RSVP) • RSVP++ • RSVP-TE, including: • Bypass LSP static configuration • Ingress LSP statistics in a file • RSVP-TE hitless-MBB with no artificial delays • 32,000 transit LSPs • Auto bandwidth • Class-based forwarding (CBF) with 16 classes • CBF with next-hop resolution • Convergence and scalability • Graceful restart signaling • JTI interface statistics and LSP event export • LSP next-hop policy

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• LSP self-ping • MPLS fast reroute (FRR) • MTU signaling • Optimize adaptive teardown • Node/link protection • Refresh reduction • Soft preemption • Shared Risk Link Group (SRLG) • Static LSPs with IPv4 nexthop, IPv6 next-hop, and IPv6 nexthop with next-table support for bypass • Traffic engineering, including: • TE++: Dynamic ingress LSP splitting • Traffic engineering extensions (OSPF-TE and ISIS-TE) • Traffic engineering options: `bgp`, `bgp-igp`, `bgp-igp-both-ribs`, and `mpls-forwarding` [See MPLS Applications User Guide .] • Support for an increased scale of transit RSVP-TE–signaled MPLS label-switched paths (LSPs) that are enabled with link protection. • Enhanced scaling for the following MPLS features: • RSVP transit LSPs with link and node protection • RSVP ingress and egress LSPs with ultimate-hop popping (UHP) and penultimate-hop popping (PHP) • LDP-over-RSVP LSPs

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Packet Forwarding Engine statistics • Fast reroute (FRR) and make before break (MBB) • Weighted ECMP • Ping and traceroute • Clone route • Transit statistics [See MPLS Applications User Guide .] • Support for RSVP-based and LDP-based point-to-multipoint (P2MP) LSPs with graceful restart. In addition, the router supports IP unicast traffic in a label-edge router (LER) role and both IP unicast and multicast traffic in a label-switching router (LSR) role. [See Point-to-Multipoint LSPs Overview.] • Support for MPLS features P2MP ping and P2MP LSPs traceroute. MPLS ping and traceroute provide the mechanism to detect data-plane failure and isolate faults in the MPLS network. The traceroute or ping is initiated to validate LSP paths on P2MP. [See MPLS Applications User Guide .] • Optimized fast branch updates. The method of making fast-branch updates to a multicast replication tree has been refined. Now, any membership changes in the tree trigger fast make-before-break (FMBB) re-optimization of the tree and ensure that there is no traffic loss. [See Multicast Shortest-Path Tree.]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Multicast	• MVPN BIER with MPLS encapsulation. Junos OS Evolved supports the Bit Index Explicit Replication (BIER) architecture to simplify control and forwarding planes by eliminating the need for multicast trees and per-flow states. With BGP-MVPN as an overlay, you can configure BIER-enabled provider tunnels for multicast VPNs. [See BIER Overview and bier.] • IS-IS as routing underlay for BIER. Junos OS Evolved supports the advertisement of BIER information of one or more BIER sub-domains using IS-IS as the IGP underlay. Key BIER information such as BFR IDs and BFR prefixes in each sub-domain are flooded through the IS-IS domain to generate the BIER forwarding table. [See IS-IS Extension for BIER and bier-sub-domain (Protocols IS-IS).] • IPv4 and IPv6 multicast support including MSDP, support for PIM-SM as the first-hop router (FHR) or last-hop router (LHR), and for anycast, static, or local rendezvous point (RP). • Support for multicast-only fast reroute (MoFRR) for both IPv4 and IPv6 traffic flows. MoFRR minimizes multicast packet loss in PIM domains when there are link failures. MoFRR is supported for PIM sparse mode (SM) and source-specific multicast (SSM) modes only. Support does not extend to Multipoint LDP-based MoFRR. [See **Understanding Multicast-Only Fast Reroute**.] • Support for bidirectional Protocol Independent Multicast for multicast traffic. [See **pim-snooping**.]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Routing policy and firewall filters	• **Support added to hierarchical policers for applying user-selectable bandwidth for premium and non-premium traffic**— Use the firewall filter action `policer-charge` to subtract available bandwidth credits and make bandwidth available to the aggregate policer. • Firewall output filtering support using Fast Lookup Filter (FFT) block for line-rate performance of up to 2 billion PPS. The `fast-lookup-filter` statement from the CLI filter configuration prioritizes output filtering (but not input filtering) on the FFT block. FFT enables support for 128 unique output filters across IPv4, IPv6, or MPLS families. [See fast-lookup-filter (PTX).] • SNMP MIB support for `jnxFirewallCounterTable` object. Junos OS Evolved SNMP extends support for the `jnxFirewallCounterTable` and its objects: • `jnxFirewallCounterEntry` • `jnxFWCounterPacketCount` • `jnxFWCounterByteCount` • `jnxFWCounterDisplayFilterName` • `jnxFWCounterDisplayName` • `jnxFWCounterDisplayType` [See SNMP MIB Explorer.] • Firewall filters support. IPv4 and IPv6 firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on the router from a source address to a destination address. [See Firewall Filter Match Conditions and Actions (PTX Series Routers).] • Support for filter-based forwarding.

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	[See Example: Configuring Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address.] • Support for SDN-based network to configure certain router interfaces to pass traffic toward an SDN controller. Use firewall filters to match and redirect packets defined at the `[edit services inline-monitoring instance]` hierarchy level. Supported match criteria includes IPv4, IPv6, and family `any` (destination), VLAN ID, and certain traceroute redirect packets. [See controller.] • Support for firewall filters on discard interfaces. You can apply firewall filters on a discard interface. The action specified by the filter (log or count) is executed before the traffic is discarded. Firewall filters are supported only for IPv4 and IPv6 traffic in the egress direction of the interface. [See Configuring Firewall Filters.] • **Support for firewall features, including:** • Forwarding IPv4 and IPv6 • Firewall filter • Load balancing • MPLS fast reroute • Host path • Egress peer engineering [See Firewall Filter Match Conditions and Actions (PTX Series Routers).] • Support for input-chain and output-chain CLI filters. Use multiple levels of CLI filters. The filter chain helps in logically grouping filters with a specific pattern of rules, instead of evaluating all the filter terms in one filter and deciding at the last term of it. The feature provides you flexibility in modelling

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	the filters as and when it is applicable in the solution. You can configure up to eight filters in both input chains and output chains. [See Example: Using Firewall Filter Chains, output-chain, and input-chain.] • Support for nested filters, which enable you to reference a common firewall filter by attaching it to multiple firewall policies (a filter being one or more match conditions and corresponding actions). You can bind nested filters to the following interface types: • `inet`—Both input and output directions • `inet6`—Both input and output directions • `mpls`—Input direction only You can also bind the filters to routing instances, and in the input direction, in the output direction, or in both directions. [See Guidelines for Nesting References to Multiple Firewall Filters and Example: Nesting References to Multiple Firewall Filters.] • Support for matching `ip-options` in IPv4 packet headers. Use the `ip-options any` match condition to match fields in the IPv4 header and create firewall filter rules to handle the matched packets. Specifying `ip-options` provides a finer level of control, so for example, you can create a rule to drop any IPv4 packets that do not include at least one IP option in the header. Configure the match condition at the `[edit firewall family inet filter` *`name`* `term` *`name`* `from ip-options any]` hierarchy level. [See Firewall Filter Match Conditions for IPv4 Traffic .] • Support for labeling interfaces with specified group IDs from 1 to 255 and matching interface-group ID on the firewall filter. The filter recognizes which interface the packet comes from and performs actions only specified for a certain interface group.

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	[See Understanding BGP Flow Routes for Traffic Filtering .] • Firewall filter support for bitwise logical operations for TCP flag match. [See Firewall Filter Match Conditions Based on Bit-Field Values .] • MPLS filter payload match. IPv4 and IPv6 payload fields match conditions are available for MPLS traffic. Additionally, the following match conditions are available: • MPLS header EXP match conditions for MPLS traffic—`exp0`, `exp1`, `exp0-except`, `exp1-except`. Existing match conditions `exp` and `exp-except` will be deprecated. • MPLS header Label match conditions for MPLS traffic—`label0`, `label1`, `label0-except`, `label1-except`. Existing match conditions `label` and `label-except` will be deprecated. • MPLS header TTL match conditions for MPLS traffic—`ttl0`, `ttl1`, `ttl0-except`, `ttl1-except`. Existing match conditions `ttl` and `ttl-except` will be deprecated. • MPLS header Bottom of Stack match conditions for MPLS traffic—`bottom-of-stack0` and `bottom-of-stack1` [See Firewall Filter Match Conditions for MPLS Traffic .] • Unicast RPF support for both IPv4 and IPv6 traffic flows. [See Configuring Unicast RPF Loose Mode.] • Enhanced scaling for DoS and protection offers loose mode unicast RPF on IPv4 and IPv6. [See Configuring Unicast RPF Loose Mode.] • Support for DCU and SCU accounting. Source class usage (SCU) accounting provides breakdown of output interface traffic statistics that originates from specific prefixes. Destination class usage (DCU) accounting provides breakdown

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	of input interface traffic statistics that is destined for specific prefixes. [See Understanding Source Class Usage and Destination Class Usage Options.] • Class-based firewall filters. You can apply firewall filters actions such as `drop`, `reject`, `sample`, and `police` on packets classified by destination class usage (DCU) and source class usage (SCU) accounting. You can use this feature, for example, as part of a design to provide distributed denial-of-service () protection to specific customers. [See Configure the Filter Profile.] • Support for forwarding class and packet loss priority as policer actions. You can use forwarding class (FC), and both FC and packet loss priority (PLP) together, as policer actions in policer policy configurations. This includes both ingress and egress directions. • Support for 2-color Layer 3 interface policers (ingress and egress). [See Basic Two-Rate Three-Color Policers.] • Support for packet-rate policers. You can use a count of packets as the threshold for traffic policers. Per-packet policers can better mitigate low-and-slow types of denial-of-service (DoS) and distributed denial-of-service () attacks. You can apply packet-level policers in the ingress or egress interface direction. These policers support both two-color and three-color policies. The following families are supported: `inet`, `inet6`, `mpls`, and `ethernet-switching` . Configure per-packet policer rates at the `[edit firewall policer` *`policer-name`*] hierarchy level using the `pps-limit` (packets per second) and `packet-burst-size-limit` (packets) configuration statements. [See Packets-Per-Second (pps)-Based Policer Overview and pps-limit (Policer).]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Shared-bandwidth and percentage policer. Use the shared-bandwidth policer for instances where policers are attached to aggregated Ethernet interface bundles with child legs spanning different Packet Forwarding Engine or Flexible Port Concentrator (FPC) instances. The bandwidth policers program the policer token bucket with weighted bandwidth or burst (depending on the number of child legs per Packet Forwarding Engine). The percentage policer feature allows you to configure the bandwidth policer relative to the physical-interface speed where you configure the class-of-service (CoS) shaping rate. After the configuration, the egress policer can then use this base CoS shaping rate instead of the physical-interface speed. [See Configure the Filter Profile.] • Two-color and three-color traffic policers for input and output traffic. The supported actions are `discard`, `forwarding-class`, and `loss-priority` (`high` and `low`). You can attach policers to logical interfaces and the protocol families `mpls`, `inet`, and `inet6`. [See Basic Two-Rate Three-Color Policers.] • Filter-based GRE encapsulation and de-encapsulation and filter-based MPLS-in-UDP de-encapsulation. We've enabled the following encapsulation and de-encapsulation workflow: 1. An incoming packet matches a filter term with an encapsulate action. The packet is encapsulated in an IP +GRE header and is forwarded to the endpoint's destination. `set firewall tunnel-end-point` *`tunnel-name`* `ipv4|ipv6 source-address` *`address`* `set firewall tunnel-end-point` *`tunnel-name`* `ipv4|ipv6 destination-address` *`address`* `set firewall tunnel-end-point` *`tunnel-name`* `gre` `set firewall family inet|inet6 filter` *`name`* `term` *`name`* `from source-address` *`address`*

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	`` set firewall family inet|inet6 filter name term name then encapsulate tunnel-name set firewall family inet|inet6 filter name term last then accept set interfaces interface-name unit number family inet|inet6 filter input set interfaces interface-name unit number family inet|inet6 address address # This source address differs from the one for the tunnel endpoint. ``` 2. At the destination, the packet matches a filter term with a de-encapsulate action. The GRE header or MPLS-in-UDP header is stripped from the packet. The inner packet is routed to its destination. ``` set firewall family inet|inet6 filter name term name from source-address address set firewall family inet|inet6 filter name term name from protocol gre set firewall family inet|inet6 filter name term name then decapsulate gre # Optionally de-encapsulate mpls-in-udp. set firewall family inet|inet6 filter name term last then accept set interfaces interface-name unit number family inet|inet6 filter input filter-name set interfaces interface-name unit number family inet|inet6 address address # This is the destination address. ``` [See Components of Filter-Based Tunneling Across IPv4 Networks and tunnel-end-point.] • Support for tunnel de-encapsulation using firewall filters for GRE and UDP tunnels. [See Configuring a Filter to De-Encapsulate GRE Traffic and decapsulate (Firewall Filter).]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Routing protocols	• BGP flow specification. BGP can carry flow-specification network layer reachability information (NLRI) messages on PTX10002-36QDD devices with 14.4 Tbps line cards. Propagating firewall filter information as part of BGP enables you to propagate firewall filters against denial-of-service (DOS) attacks dynamically across autonomous systems. The following match conditions are not supported: • ICMP codes alone [`inet/inet6`] • Source/destination prefix with offset for `inet6` • Flow label for `inet6` fragment [for `inet6`] Junos OS Evolved running on this router doesn't support the traffic marking action. To configure flow routes statically, configure the match conditions and actions at the `[edit routing-options]` hierarchy level. • Forwarding IPv6 transit statistics. [See **BGP User Guide**.]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Network management and monitoring	• Local port mirroring support. You can use port mirroring to copy packets entering or exiting a port or entering a VLAN and to send the copies to a local interface for local monitoring. The following features are included: • Interface filter on ingress and egress • Forwarding table filter (FTF) on ingress • Families `inet` and `inet6` • Aggregated Ethernet interfaces at both ingress and egress Use the following CLI hierarchies to configure port mirroring: • `[edit interfaces]` • `[edit forwarding-options port-mirroring]` • `[edit firewall filter]` You can configure family `inet` and family `inet6` in the `[edit interfaces]` and the `[edit forwarding-options port-mirroring]` hierarchies for this feature. This feature applies to global port mirroring only. [See Understanding Port Mirroring and Analyzers.] • Remote port mirroring with ToS or DSCP settings. You can send sampled copies of incoming packets to remotely connected network management software. You send the packets using GRE, which is supported by flexible tunnel interfaces (FTIs). You can set ToS and DSCP values to provide necessary priorities in the network for these packets. You can also apply policing to sampled packets that are leaving the FTI. Configure the settings you need in the `[edit forwarding-options port-mirroring instance` *`instance-name`* `output]` hierarchy. [See instance (Port Mirroring).]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Support for additional family `any` in port mirroring You can configure family `any` (as well as the earlier family options, `inet` and `inet6`) for local port mirroring and remote port mirroring. You can use the family `any` configuration option to process the families `any`, `ccc`, `ethernet-switching`, or `mpls`. **NOTE**: You use the family `any` configuration option to process all 4 families. Use `[edit forwarding-options port-mirroring]` for local port mirroring or `[edit forwarding-options port-mirroring instance` *`instance-name`*`]` for remote port mirroring, with both of those configurations also requiring a firewall filter. The following configuration statements are no longer part of the port mirroring configuration on PTX Series devices: • `next-hop` for family `any` • family `vpls` • `no-filter-check` • `hosted-service` • `server-profile` [See port-mirroring.] • Support for EVPN-VXLAN filtering and port mirroring based on VNI match conditions. You can construct a firewall filter to filter EVPN-VXLAN traffic by using the VXLAN network identifier (VNI) values in the match condition on ingress and egress interfaces. This feature supports redirecting traffic to a global port-mirroring instance. To filter traffic based on the VNI, use the following commands: `set firewall filter` *`filter-name`* `term` *`term-name`* `from vxlan vni` *`vni-value`*

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	`set firewall filter` *`filter-name`* `term` *`term-name`* `from vxlan vni-except` *`vni-value`* `vni-value` can be a numeric value or range of numeric values. [See Firewall Filter Match Conditions and Actions (PTX Series Routers).] • Support for the sFlow technology, which is a monitoring technology for high-speed switched or routed networks. The sFlow monitoring technology randomly samples network packets and sends the samples to a monitoring station. [See Understanding How to Use sFlow Technology for Network Monitoring.] • sFlow support for MPLS interfaces to sample and report MPLS traffic on the routers. [See sFlow Technology Overview.] • Ingress and egress sFlow functionalities for transit nodes are supported for IPv4-in-IPv4, IPv6-in-IPv4 and regular IPv4/IPv6 traffic. In transit-only devices, the IP-in-IP encapsulated packet can transit through the device without any change or might get de-encapsulated and forwarded or de-encapsulated and encapsulated and forwarded based on the next-hop configuration. Additionally, the packet might traverse through multiple VRFs while getting forwarded. The router supports ingress and egress sFlow for all those variations. [See sFlow Monitoring Technology.] • sFlow technology support for exporting extended IPv4 and IPv6 tunnel egress structure. sFlow technology supports the export of the Extended Tunnel Egress Structure fields for traffic entering IPv4 or IPv6 GRE tunnels. These additional attributes provide information about the GRE tunnel into which a packet entering the device will get encapsulated. The GRE tunnel could be IPv4 or IPv6. The feature is supported only when sFlow is enabled in the ingress direction wherein firewall-based GRE happens on IPv4 or IPv6 packets.

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	The device supports the feature for the following traffic scenarios when ingress sFlow sampling is enabled: • Incoming IPv4 traffic that undergoes IPv4 GRE • Incoming IPv6 traffic that undergoes IPv4 GRE • Incoming IPv4 traffic that undergoes IPv6 GRE • Incoming IPv6 traffic that undergoes IPv6 GRE [See sFlow Monitoring Technology.] • Sample size support in sFlow. 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 sFlow Monitoring Technology.] • Support for passive monitoring, including support for passive monitoring on MPLS-encapsulated packets. You can configure passive monitoring on any interface on the PTX Series routers, and you can use this feature to monitor MPLS-encapsulated packets. After you enable passive monitoring, the router accepts and monitors traffic on the interface and forwards those packets to monitoring tools such as IDS servers and packet analyzers, or to other devices such as other routers or end-node hosts. [See Passive Monitoring and passive-monitor-mode.] • **Support for link fault management (LFM)**—We support IEEE 802.3ah OAM LFM to monitor point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The following LFM features are supported: • Link discovery with active and passive modes • Detect-LOC • Remote loopback

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Loopback tracking • Action profile • GRES and non-graceful Routing Engine switchover [See Introduction to OAM Link Fault Management (LFM).]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Segment routing	• Segment routing support. You can configure the following Source Packet Routing in Networking (SPRING) or segment routing features on the router: • MPLS (segment routing using IS-IS): • Ping and traceroute for single IS-IS node or prefix segment • BGP Link State (BGP-LS): • Segment routing extensions for IS-IS • Segment routing extensions for OSPF • BGP: • Binding segment identifier (SID) for segment routing–traffic engineering (SR-TE) • Binding SID for SR-TE [draft-previdi-idr-segment-routing-te-policy] • Programmable routing protocol process APIs for SR-TE policy provisioning • Static SR-TE policy with mandatory color specification • Static SR-TE policy without color specification • IS-IS: • Adjacency SID • Advertising maximum link bandwidth and administrative color without RSVP-TE configuration • Anycast and prefix SIDs • Configurable segment routing global block (SRGB)

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Node and link SIDs • Segment Routing Mapping Server (SRMS) and client • Topology-independent loop-free alternate (TI-LFA): • Link and node protection for IPv4 addressing (not required for IPv6 prefixes) • Link and node protection for IPv4 addressing (required for IPv6 prefixes) • Protection for SRMS prefixes • OSPF: • Advertising maximum-link bandwidth and administrative color without RSVP-TE configuration • Anycast SID • Configurable SRGB • Inter-area support • Node and link SID • Prefix SID • Segment Routing Mapping Server (SRMS) and client • Static adjacency SID • TI-LFA: • Link and node protection • Protection for SRMS prefixes

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• MPLS ping and traceroute for single OSPF node or prefix segment • IGP adjacency SID hold time • Path Computation Element Protocol (PCEP) for segment routing LSPs • • BGP IPv4 labeled-unicast resolution over: • BGP IPv4 SR-TE with IPv4 segment routing using IS-IS and OSPF • Non-colored IPv4 SR-TE with segment routing using IS-IS and OSPF • Static colored IPv4 SR-TE with segment routing using IS-IS and OSPF • • BGP Layer 3 VPN over: • Colored SR-TE tunnels and IPv4 protocol next hops • Non-colored SR-TE tunnels and IPv4 protocol next hops • BGP-triggered dynamic SR-TE colored tunnels • Class-based forwarding and forwarding table policy LSP next-hop selection among non-colored SR-TE LSPs • First-hop label support for SID instead of an IP address • Path specification using router IP addresses (segment routing segment list path ERO support using IP address as next hop and loose mode) • SR-TE color mode: • 00—Route resolution fallback to IGP path

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• 01—Route resolution fallback to color only null routes • Static LSPs with member-link next hops for aggregated Ethernet bundles (also known as adjacent SID per LAG bundle or aggregated Ethernet member link) [See Understanding Source Packet Routing in Networking (SPRING).] • Support for scaled-up static and BGP segment routing policies, where each policy contains eight segment routing paths with five labels per path without make-before-break (MBB). [See egress-chaining and fib-next-hop-split .] • SPRING statistics sensor support for JTI supports export of SPRING statistics to an outside collector by using remote procedure call (gRPC) services. The feature provides the segment-identifier (SID)-level and interface-level traffic counts for Source Packet Routing in Networking (SPRING) traffic. These statistics reflect the SPRING LSP utilization in the traffic engineering database, which aids in correctly rerouting the RSVP LSPs. To enable SPRING statistics, include the following statements on the client device: • For egress (per-interface egress), use `set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link egress` • For egress (per-SID egress), use `set protocols isis source-packet-routing sensor-based-stats per-sid egress` • For ingress (per-SID ingress), use `set protocols isis source-packet-routing sensor-based-stats per-sid ingress`. Use the following sensors to export statistics by means of gRPC services to an outside collector:

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• **/junos/services/segment-routing/interface/egress/usage/** for egress (per-interface egress) aggregate SPRING traffic. • **/junos/services/segment-routing/sid/usage/** for egress (per-SID egress) and ingress (per-SID ingress) aggregate SPRING traffic. [See source-packet-routing and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface.] • BGP and statically configured SR-TE traffic statistics sensor support for JTI. [See source-packet-routing, Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface, and Understanding OpenConfig and gRPC on Junos Telemetry Interface.] • Support for segment routing over UDP. Configure the `udp-tunneling encapsulation` statements at the `[edit protocols isis source-packet-routing]` hierarchy level to enable SR-MPLS routers and IP-only routers to seamlessly coexist, by encapsulating SR-MPLS label stacks in IP/UDP encapsulation. This feature also supports: • Entropy in the UDP source port • Underlay and overlay ECMP at the start of the tunnel • Policy control to resolve dynamic tunnels SR-over-UDP supports tunnels without a loopback stream in the Packet Forwarding Engine, thereby reducing additional bandwidth consumption. [See Next-Hop-Based Dynamic Tunnels and source-packet-routing (Protocols IS-IS).] • SPRING : JTI : ingress SR-TE statistics per binding SID and segment list (static, BGP, PCEP paths). Use this feature to provide route statistics for segment routing-traffic engineering (SR-TE) per label-switched path (LSP). Junos OS Evolved uses

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	Junos telemetry interface (JTI) and gRPC services to provide the statistics. Supported resource paths (sensors) include: • **/junos/services/segment-routing/traffic-engineering/tunnel/lsp/ingress/usage/** • **/junos/services/segment-routing/traffic-engineering/tunnel/lsp/transit/usage/** [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface, source-packet-routing, and show spring-traffic-engineering.]
	• SPRING statistics sensor support for JTI supports export of Source Packet Routing in Networking (SPRING) statistics to an outside collector by using remote procedure call (gRPC) services and gRPC Network Management Interface (gNMI) services. This feature provides interface-level and segment identifier (SID)-level ingress statistics. The feature also provides egress statistics for each child member at the physical interface level. To enable SPRING statistics, include the following statements on the client device: • For egress (per-child member at the physical- interface level), use the `set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link egress` command at the `[edit]` hierarchy level. • For ingress (per-SID ingress and per-interface ingress), use the `set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link ingress` command at the `[edit]` hierarchy level. Use the following sensors to export statistics by means of gRPC or gNMI services to an outside collector:

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• **/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/in-octets/** for ingress (per-SID ingress and per-interface ingress) SPRING traffic. • **/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/out-octets/** for egress (per-child member at the physical-interface level) SPRING traffic. • **/network-instances/network-instance/mpls/signaling-protocols/segment-routing/interfaces/interface/state/out-pkts/** for egress (per-child member at the physical-interface level) SPRING traffic. [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface and source-packet-routing .] • Support for segment-routing telemetry sensor enhancements. We support segment routing sensor enhancements for SID-level and interface-level traffic counts. These enhancements comply with the current supported sensors in the OpenConfig models **openconfig-segment-routing.yang** and **openconfig-mpls.yang**. • SR-TE colored policy RIB5 and SR-TE colored telemetry sensor support. We support JTI streaming and ON-CHANGE sensors that deliver operational state statistics for Segment Routing–Traffic Engineering (SR-TE) colored policy RIB5 and SR-TE colored telemetry sensors. Statistics are delivered to an outside collector using gRPC or gNMI. The feature includes new OpenConfig resource paths for existing and new SR-TE policy (tunnel) and SR-TE per-LSP colored statistics. [See Telemetry Sensor Explorer .] • Support for SRv6 network programming in IS-IS. Use this feature to configure segment routing in a core IPv6 network without an MPLS dataplane.

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	To enable SRv6 network programming in an IPv6 domain, include the `srv6` statement at the `[edit protocols isis source-packet-routing]` hierarchy level. To advertise the Segment Routing Header (SRH) locator with a mapped flexible algorithm, include the `algorithm` statement at the `[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 SRv6 network programming and Layer 3 Services over SRv6 in BGP. You can configure BGP-based Layer 3 service over an SRv6 core. You can enable Layer 3 overlay services with BGP as the control plane and SRv6 as the data plane. SRv6 network programming provides flexibility to leverage segment routing without deploying MPLS. Such networks depend only on the IPv6 headers and header extensions for transmitting data. To configure IPv4 and IPv6 transport over an SRv6 core, include the `end-dt4-sid` *`sid`* and the `end-dt6-sid` *`sid`* statements at the `[edit protocols bgp source-packet-routing srv6 locator name]` hierarchy level. To configure IPv4 VPN and IPv6 VPN service over an SRv6 core, include the `end-dt4-sid` *`sid`* and the `end-dt6-sid` *`sid`* statements at the `[edit routing-instances` *`routing-instance-name`* `protocols bgp source-packet-routing srv6 locator` *`name`*`]` hierarchy level. [See Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP .]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• **Operations, Administration and Management (OAM) ping support for segment routing with IPv6 (SRv6) network programming.** You can perform an OAM ping operation for any SRv6 segment identifier (SID) whose behavior allows upper layer header processing for an applicable OAM payload. As segment routing with IPv6 data plane (SRv6) adds only the new type-4 routing extension header, you can use the existing ICMPv6-based ping mechanisms for an SRv6 network to provide OAM support for SRv6. Ping with O-Flag (segment header) is not supported. [See ITU-T Y.1731 Ethernet Service OAM Overview and How to Enable SRv6 Network Programming in IS-IS Networks .] • Support for SRv6 traceroute. We support the traceroute mechanism for segment routing for IPv6 (SRv6) segment identifiers. You can use traceroute for both UDP and ICMP probes. By default, traceroute uses UDP probes. For ICMP probes, use the `traceroute` command with the `probe-icmp` option. [See How to Enable SRv6 Network Programming in IS-IS Networks .] • SRv6 support for static SR-TE policy. You can configure static segment routing–traffic engineering (SR-TE) tunnels over an SRv6 data plane. Use the following configuration commands to enable SRv6 support: • For an SR-TE policy: `set protocols source-packet-routing srv6` • For an SR-TE tunnel: `set protocols source-packet-routing source-routing-path lsp` *`name`* `srv6` • For an SR-TE segment list: `set protocols source-packet-routing source-routing-path segment-list srv6` [See Understanding SR-TE Policy for SRv6 Tunnel .]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
Services applications	• Inline active flow monitoring support. [See Understand Inline Active Flow Monitoring .] • Inline monitoring services support for packet mirroring with metadata. [See Inline Monitoring Services Configuration .] • Support for additional RPCs for the gNOI certificate management (`cert`) service. Junos OS Evolved supports the following gRPC Network Operations Interface (gNOI) `cert` 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 .] • CFM support: • Up maintenance association end points (MEPs) in distributed periodic packet management (PPM) • Distributed Y.1731 on synthetic loss measurement (SLM), delay measurement (DM), and loss measurement (LM) • Down MEPs on bridges, circuit cross-connect (CCC) , and Ethernet VPN (EVPN) • Distributed session support for connectivity fault management (CFM) on aggregated Ethernet • Enhanced CFM mode • IPv4 (`inet`) support for Data Model (DM) and synthetic loss message (SLM)

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Action profile for marking a link down, except for EVPN and bridge up MEP • LM colorless mode • DM and LM on aggregated Ethernet if all active child links are on the same Packet Forwarding Engine • Supported CFM protocol data units (PDUs), as follows: • Continuity check messages (CCM) • LBM • LBR • Link Trace Message (LTM) • Link Trace Reply (LTR) • 1DM • Delay measurement message (DMM) • Delay measurement reply (DMR) • LMM • LMR • Synthetic loss message (SLM) • Synthetic loss reply (SLR) • Enterprise and service provider configurations • VLAN normalization • VLAN transparency for CFM PDUs • CoS forwarding class (FC) and CoS packet loss priority (PLP) for CFM

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• CFM session on child physical interface in distributed mode • SNMP • Chassis ID or Send ID type, length, and value • Trunk mode • Maintenance association intermediate point (MIP) [See Connectivity Fault Management (CFM).] • Support for enhanced connectivity fault management.extends CFM support to inline mode. Support includes: • Up and down maintenance association end points (MEPs) on bridges, circuit cross-connect (CCC), and Ethernet VPN (EVPN) in inline mode • ITU-T Y.1731 on synthetic loss measurement (SLM) and delay measurement (DM) • Inline session support for connectivity fault management (CFM) on aggregated Ethernet • Enhanced CFM mode by default • Supported inline performance monitoring (PM) sessions, as follows: • PM Tx • PM Rx • PM responder • IPv4 (`inet`) and IPv6 (`inet6`) support for continuity check messages (CCM), delay measurement (DM), and synthetic loss message (SLM) • DM on aggregated Ethernet with at least one child link on the anchor Packet Forwarding Engine

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Action profile for marking a link down, except for EVPN and bridge up MEP • Supported CFM protocol data units (PDUs) for inline handling, as follows: • CCM • Delay measurement message (DMM) • Delay measurement reply (DMR) • Synthetic loss message (SLM) • Synthetic loss reply (SLR) • Enterprise and service provider configurations • VLAN normalization • VLAN transparency for CFM PDUs • Combination of up MEP, down MEP, or maintenance association intermediate point (MIP) configuration over the same interface [See Connectivity Fault Management (CFM).]
Security services	• Support for DDoS IS-IS classification and higher DDoS bandwidth for Layer 2 and Layer 3 protocols. [See **show ddos-protection protocols isis** and **protocols (DDoS) (ACX Series, PTX Series, and QFX Series)**.]
Software installation and upgrade	• Support for secure BIOS and secure boot implementation based on the UEFI 2.4 standard. [See Secure Boot.]

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
VPNs	• MPLS-based Layer 3 VPNs support includes: • MPLS over Layer 3 VLAN-tagged subinterfaces • Per-next-hop label allocation • Mapping of the label-switched interface (LSI) logical interface label to the VPN routing and forwarding (VRF) routing table using the `vrf-table-label` statement • ICMP tunneling and MPLS traceroute • Disabling time-to-live (TTL) decrementing using `no-propagate-ttl` [See Layer 3 VPNs Feature Guide for Routing Devices.] • Carriers-of-carriers and inter-AS VPN supported features include: • Carrier-of-carriers VPN service • Interprovider Layer 3 VPN Option A • Interprovider Layer 3 VPN Option B • Interprovider Layer 3 VPN Option C However, traffic statistic collection for BGP labeled unicast is not supported for carrier-of-carrier VPNs and interprovider traffic. [See Carrier-of-Carrier VPNs.] • Layer 2 VPN feature support includes: • Transport of Layer 2 frames over MPLS (LDP signaling) • Layer 2 VPNs over tunnels (BGP signaling) • Simple Ethernet and VLAN-based cross-connect (also known as connections)

Table 1: Features Supported on the PTX10002-36QDD *(Continued)*

Feature	Description
	• Local and remote switching • Ethernet and VLAN CCC • Single-tagged CCC logical interfaces • Control word • Regular and aggregated Ethernet interfaces • Layer 2 protocol pass-through • Layer 2 circuit backup interface and backup neighbor • Layer 2 circuit statistics and CoS • VCCV with type 2 and type 3 [See Layer 2 VPNs and VPLS User Guide for Routing Devices and TCC Overview.]

Chassis

- **Support for powering on, powering off, or restarting Packet Forwarding Engine (PTX10002-36QDD)**—Starting in Junos OS Evolved Release 23.4R2-S1, you can power off, power on, or restart the Packet Forwarding Engines in the router by following these steps:

  1. Configure the *pair* of Packet Forwarding Engines that you want to restart or power off/on—for example:
     - `set chassis fpc 0 pfe 0 power on`
     - `set chassis fpc 0 pfe 1 power on`

     NOTE: The four Packet Forwarding Engines are numbered 0–3. You configure them in pairs—0 and 1; 2 and 3.

  2. Issue the power on, power off, or restart command—for example:
     - `request chassis fpc slot 0 pfe 0 power on`
  3. Enter **yes** when the following question appears on the screen:

- ```
  Warning: pfe 1 will also be offlined. Do you wish to continue?
          [yes,no]
  ```

> **NOTE**: You can also set the `reset-pfe` action to reset a PFE when a chassis error occurs—configure the `action` statement at `[edit chassis fpc slot-number error error-severity-level]` hierarchy level.

[See request chassis fpc and action (chassis error).]

- **Power redundancy and resiliency support (PTX10016)**–Starting in Junos OS Evolved Release 24.2R2, the PTX10016 with the JNP10K-PWR-AC3 power supply modules (PSMs) supports the following features:

  - N+1 power redundancy
  - Resiliency support for FRU components
  - PSM watchdog

  You can enable either source redundancy or feed redundancy for the PSM.

  [See Managing Power, watchdog (PSM), and Chassis-Level User Guide.]

Class of Service

- **Policy map support (PTX10002-36QDD)**—Starting in Junos OS Evolved Release 23.4R2-S1, PTX10002-36QDD routers support policy maps, which enable you to assign rewrite rules on a per-customer basis. The policy map makes it possible to use any packet field to identify a given flow and specify a rewrite value for that flow. PTX10002-36QDD routers support the following types of packet marking: INET-Precedence, DSCP, IEEE802.1p, and IEEE802.1ad.

  You can define a `policy-map` at the `[edit class-of-service]` hierarchy level. You enable `policy-map-marking` at the egress interface at the `[edit class-of-service interfaces interface-name unit unit-number]` hierarchy level.

  [See Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps.]

MACsec

- **MACsec bounded delay protection (PTX10002-36QDD)**—Starting in Junos OS Evolved Release 23.4R2-S1, you can enable Media Access Control Security (MACsec) bounded delay protection to protect your network against man-in-the-middle attacks. MACsec is an industry-standard security technology capable of identifying and preventing most security threats. During a man-in-the-middle
  ```

attack, an attacker intercepts packets and might redirect or modify them. This attack can cause an unexpected delay in how long a packet or frame takes to arrive at its intended destination.

When you enable MACsec bounded delay protection, the device guarantees that a frame will not be delivered after a delay of two seconds or more. MACsec periodically compares the number of frames transmitted to the number received. If a frame is sent but not received within two seconds, MACsec drops the packet. This ensures that a delay of MACsec frames resulting from a man-in-the-middle attack will not go undetected.

To enable bounded delay protection, configure the following options at the `[edit security macsec connectivity-association connectivity-association-name]` hierarchy level:

- `mka bounded-delay`
- `replay-protect replay-window-size 0`

[See Configuring Bounded Delay Protection.]

Precision Time Protocol (PTP)

- **Support for G.8275.1 profile, PTP over Ethernet encapsulation, and hybrid mode over LAG with PTP over Ethernet (PTX10002-36QDD)**—Starting in Junos OS Evolved Release 23.4R2-S1, these features are added to the router.

  [See G.8275.1 Telecom Profile, Guidelines for Configuring PTP over Ethernet, andHybrid Mode.]

Services Applications

- **Inline active flow monitoring IPFIX and version 9 template support for CoS policy-map name reporting in the ingress direction (PTX10002-36-QDD)**—Starting in Junos OS Evolved Release 23.4R2-S1, we support a new Juniper-specific enterprise Information Element ID, 32765, in the data record templates `ip4-template` and `ipv6-template`. This new IE ID is 4 bytes long and contains the first 4 characters of the policy-map name. Therefore, the first four letters of your policy-map names should be unique. You configure this new IE ID with the `include-policy-map-name` statement at the `[edit services flow-monitoring (version-ipfix|version9) <template-name> data-record-fields]` hierarchy level. You configure policy maps at the `[edit class-of-service policy-map]` hierarchy level.

  [See data-record-fields, Understand Inline Active Flow Monitoring, and Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps.

Software Installation and Upgrade

- **Zero touch provisioning on WAN interfaces (PTX10002-36QDD)**—Starting in Junos OS Evolved Release 23.4R2-S1, Zero Touch Provisioning (ZTP) dynamically detects the port speed of WAN interfaces and uses this information to create ZTP client ports with the same speed. ZTP automatically cycles through the WAN ports until it receives Dynamic Host Control Protocol (DHCP) options from the DHCP server. The device uses the DHCP options to perform the bootstrap process.

[See Zero Touch Provisioning.]

What's New in 23.4R2

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 23.4R2, click the Group by Release link. You can collapse and expand the list as needed.

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

Additional Features

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

- **Support for using VNI match condition on egress interfaces** (PTX10001-36MR, PTX10004, PTX10008, and PTX10016)

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

What's Changed

IN THIS SECTION

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

General Routing

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

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

- On PTX10004, PTX10008, and PTX10016 routers, after executing the `request node offline` command, you must wait at least 180 seconds to execute the `request chassis cb offline` command.

- **Enhanced DDoS statistics operational command (PTX Series)**—We've enhanced the aggregate DDoS statistics output field to display the aggregate statistics for BFD and DHCP protocols. The enhanced DHCP statistics output displays the collective DHCPv4 and DHCPv6 statistics for DDoS. Earlier to this release, the aggregate DDoS statistics output displayed 0 for aggregate BFD and the aggregate DHCPv4v6.

  [See show ddos-protection protocols.]

- Disable power redundancy alarms for JNP10K-PWR-DC2 PSM (PTX10008 and PTX10016)- The JNP10K-PWR-DC2 PSM supports power redundancy across two DIP switches. When all input feeds are not connected to power supplies, it triggers a chassis alarm such as `PSM 5 Input B0 and B1 Failed`. Starting in Junos OS Evolved Release 24.2R1, you can disable this chassis alarm by using the `set chassis alarm psm` *`psm number`* `input` *`input number`* `ignore` command.

  [See JNP10K-PWR-DC2 Power Supply.]

- **DDoS protection protocols statistics update (PTX Series)**—Starting in Junos OS Evolved Release 23.2R2, the `show ddos-protection protocols statistics` displays the `Max arrival rate` and `Arrival rate` output values as expected. Earlier to this release, the `Max arrival rate` and `Arrival rate` output values were displayed larger than expected.

  [See show ddos-protection protocols parameters.]

- **DDoS violation information shows incorrect default time and date (PTX Series)**—When you clear the DDoS violation state using the `clear ddos-protection protocols states` command in Junos OS Evolved, the log message displays an incorrect default time and date. However, if you bypass the recovery time while clearing the DDoS violation state, the log message displays accurately.

  See [clear ddos-protection protocols.]

- The system now checks the port number value (z) in the 'set interfaces et-x/y/z:n' configuration for a valid port range on PTX10002-36QDD. Previously, configurations with invalid port numbers were committed successfully. With this update, the system displays a UI error message and prevents committing configurations with invalid port numbers, ensuring configuration accuracy and preventing potential issues.

- Disabled CDN auto download (Junos OS Evolved) - The PKI process periodically, by default every 24 hours, polls the CDN server for the latest default trusted CA bundle and updates the list for any changes to the trusted CAs in the bundle. If there are any changes, PKI process loads them in the background. The auto download of CA certificates might generate core files. We've disabled the service of PKI query to CDN server periodically to download the latest trusted CA bundle.

- On Junos OS Evolved, password authentication for SCP based configuration archival is supported.

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the **edit protocols evpn** hierarchy level. In most use cases, you don?t need to change the default limit. If you want to change the default limit, we recommend that you don?t set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

  [See mac-ip-limit.]

EVPN

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the **edit protocols evpn** hierarchy level. In most use cases, you don?t need to change the default limit. If you want to change the default limit, we recommend that you don?t set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

  See [mac-ip-limit.]

- **Updates to syslog EVPN_DUPLICATE_MAC messages**—EVPN_DUPLICATE_MAC messages in the System log (syslog) now contain additional information to help identify the location of a duplicate MAC address in an EVPN network. These messages will include the following in addition to the duplicate MAC address:

  - The peer device, if the duplicate MAC address is from a remote VXLAN tunnel endpoint (VTEP).
  - The VLAN or virtual network identifier (VNI) value.
  - The source interface name for the corresponding local interface or multihoming Ethernet segment identifier (ESI).

  For example: `Feb 27 22:55:13 DEVICE_VTEP1_RE rpd` *`39839`*: EVPN_DUPLICATE_MAC: MAC address move detected for 00:01:02:03:04:03 within instance=evpn-vxlan on VNI=100 from 10.255.1.4 to ge-0/0/1.0.

  For more on supported syslog messages, see System Log Explorer.

Infrastructure

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

  [See **Path MTU Discovery**.]

Junos OS API and Scripting

- **<get-trace> RPC support removed (ACX Series, PTX Series, and QFX Series)**—The `show trace application` *app-name* operational command and equivalent <get-trace> RPC both emit raw trace data. Because the <get-trace> RPC does not emit XML data, we've removed support for the <get-trace> RPC for XML clients.

Network management and Monitoring

- **`get-trace` RPC support removed (ACX Series, PTX Series, and QFX Series)**—The `show trace application` *app-name* operational command and equivalent <get-trace> RPC both emit raw trace data. Because the <get-trace> RPC does not emit XML data, we've removed support for the <get-trace> RPC for XML clients.

System Management

- **Additional `Upgrade` fields for the `show system applications detail` command (ACX Series, PTX Series, and QFX Series)**—The `show system applications detail` command and corresponding RPC include additional `Upgrade` output fields. The fields provide information about notifications and actions related to various upgrade activities.

  [See show system applications (Junos OS Evolved).]

VPNs

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

  [See min-rate.]

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

  [See min-rate.]

Known Limitations

IN THIS SECTION

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.

General Routing

- Transient parity erorr **pp_0_filter_action_0_intr_pmv_eq_zero** might seen, during filter modification with traffic running. PR1778622
- EEPROM read failures for PSM are not shown on the PSM LED since the LED of the PSM is internally controlled by the PSM firmware and not by the system software. PR1770991
- During the programming of the firmware in Junos OS Evolved, all Pri/Sec/Led/Comm firmware shows as programming even if any one single firmware is getting programmed. However, the firmware upgrade will be done only on the firmware that must be programmed based on the current and available versions. PR1774769

Open Issues

IN THIS SECTION

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

- Below is the expected performance for this 21.4. Profile Freq (Hz) NoiseTransfer_0_00391_Results PASS PASS NoiseTransfer_0_00781_Results FAIL FAIL NoiseTransfer_0_01563_Results FAIL FAIL NoiseTransfer_0_03125_Results FAIL FAIL NoiseTransfer_0_06156_Results FAIL FAIL NoiseTransfer_0_12313_Results FAIL FAIL NoiseTransfer_0_24625_Results PASS PASS NoiseTransfer_0_4925_Results PASS PASS NoiseTransfer_0_985_Results PASS PASS NoiseTransfer_1_985_Results PASS PASS NoiseTransfer_3_985_Results PASS PASS NoiseTransfer_7_985_Results PASS PASS PR1624502
- On PTX10000 platforms, fault state is latched on PSM although all input and output is OK with no chassis alarm. PR1669323
- In high scaled (beyond 14000 s,g routes - 7000 ipv4 and 7000 ipv6)NGMVPN SPMSI scenarios, core might be seen on PTX10003 platforms due to memory getting exhausted. PR1708454
- On all Junos OS Evolved platforms, VMcores are seen when MACsec (Media Access Control Security) key-chains and BGP(Border Gateway Protocol) configurations are applied through Netconf. PR1732611
- When class-of-service with shaping rate is configured on Aggregate Ethernet interfaces, and the firewall policer queries the aggregate Ethernet member and not the AE interface, the shaping rate or the policy configuration does not take effect as the shaping rate is not configured in AE member. PR1735087
- The na-grpcd might crash if there are OCST, IPAFT, cliNoise, snmp polling going on in the background. It will restart automatically and should be running fine afterwards.PR1744033

- On all Junos OS Evolved based PTX platforms with physical interface/LAG (Link Aggregation Group) interface in Layer 3 mode, the untagged control traffic on Layer 3 interface gets dropped if lport value of Layer 3 interface matches with lport value of Layer 2 aggregated ethernet (AE) interface in trunk mode. PR1745528
- Both SIB PWR/STAT H/W LEDs become Unlit/OFF by Routing Engine switchover on PTX10004, PTX10008, and PTX10016 Junos OS Evolved platforms. PR1749781
- This is a day-1 issue and currently in Junos OS Evolved on PTX10008 the fan FRU is not showing in the SNMP queries. In Junos OS Evolved only FT and FTC is shown - This is in comparison MX which also show as FAN also and its status in SNMP FRU. PR1754833
- On all Junos OS Evolved PTX platforms, RIB (Routing Information Base) and FIB (Forwarding Information Base) tables are not synchronized properly, causing the P2MP (Point-to-Multipoint) LSP (label-switched-path) traffic outage when executing the CLI command `clear rsvp session`.PR1757635
- On all Junos OS Evolved PTX10008, PTX10004, and PTX10016 platforms, when a Fan-tray is removed followed by an insertion, it will lead to fan-tray failures. PR1767111
- With DHCP trace options enabled in their regression scripts, this core was seen. Issue is random not seen always. This was enabled in script for debugging in production network this will not be enabled by default hence the issue will not be seen. Its recommended to enable DHCP trace options only for debugging not otherwise.PR1771121
- Transient parity erorr "pp_0_filter_action_0_intr_pmv_eq_zero" might seen, during filter modification with traffic running. PR1778622
- This FIB limitation is roughly 4 million routes is related to PTX systems: PTX10001-36MR PTX10000-LC1201-36CD Note: For exact scale and other details please open JTAC ticket or engage Juniper account team. PR1772732
- Low hold timer BGP sessions might flap post NSR switchover at 8000 BGP sessions scale in 23.4 release. PR1781414
- On Junos OS Evolved platforms, when configuration of Decap filter action is on egress, commit error on both RE/PFE syslog messages is seen, however there is no service impact due to this issue. PR1793356
- In case of primary-only inet6 address uses for re[01]:mgmt-0 interfaces, if there is inet6 address modification, the legacy inet6 address might still reside when do "ip -6 add show dev vmb0" under OS shell. The output of CLI command `show interfaces re[01]:mgmt-0` shows the inet6 address correctly. PR1796934
- On all Junos OS Evolved platforms, time-zone info changes to default UTC after upgrade is done with restart-upgrade. PR1803511

- On all Junos OS Evolved platforms, while executing `show log messages`", error message **sysctl kern.corefile not supported** is seen which is introduced during daemon initialisation. PR1808481
- On Junos OS Evolved platforms, when license gets expired, LICENSE_EXPIRED syslog is supposed to get generated which is missing. PR1808956
- [evpn_vxlan] [evpn_instance] ptx10001-36mr :: JDI-RCT: Error message observed "RT : Mac Route install failed for rtt table index:452 Mac:0x5e000004 BD:170 ifl:4294967295 error:Generic failure" after performing disable ae member link on TOR device. PR1810348
- On Junos OS Evolved PTX platform, the Border Gateway Protocol (BGP) session flap can be seen when inline Bidirectional Forwarding Detection (BFD) is configured under routing instance without loopback interface (lo0) leading to partial traffic drop. The issue is only seen when the interface is within the routing-instance. PR1811245
- VRRP (Virtual Router Redundancy Protocol) will not work when VRRP group 0 is configured and specific Junos OS Evolved platforms PTX10001, PTX10003, PTX10004, PTX10008, and PTX10016 are working as primary.PR1816310
- Segmentation fault on grpc timer thread (might be related to keepalive) #32085 grpc issue https://github.com/grpc/grpc/issues/32085 grpc stack needs to be upgraded to 1.53 or later. PR1722414
- In Junos Evolved OS, during scaled configuration commit, Configd daemon in UI-infra is taking more time to process the commit. To optimise the processing time, most likely there will be design change required. There is no LKWR/pass instance. Also based on the configuration scale and the time taken for the commit to complete, it seems that this is a Day-1 issue.PR1701214
- JNP10K-PWR-AC3: There is a mismatch of the physical LED and shown in the SNMP LED status in case of d2d failure. PR1831436
- JNP10K-PWR-AC3: There is a mismatch of the physical LED and shown in the SNMP LED status in case of health check failure. PR1831444

Flow-based and Packet-based Processing

- A router running sampling might see the msvcs-db daemon run at 99.9% for an extended period. PR1816542

Infrastructure

- The `request routing-engine login other-routing-engine` command might fail.PR1712705

Interfaces and Chassis

- On Junos OS Evolved PTX10000 platforms, the vmcore might be seen if any component is forcefully removed from the PCI (Peripheral Component Interconnect) bus. PR1739142

MPLS

- On all Junos OS Evolved platforms, when MPLS (Multiprotocol Label Switching) statistics is configured without LSP (Label-Switched Path) configuration, the rpd process will crash and impact the routing protocols. This leads to traffic disruption due to the loss of routing information.PR1698889
- On all Junos OS Evolved platforms, a memory leak is observed in the traffic engineering database (TED) when the inet table is not cleaned up after routing instance deactivation. PR1701800

Multicast

- On vPTX platforms, the PFE (packet forwarding engine) receives an invalid token from RPD (Routing Engine daemon) for composites next-hops due to which the PFE will crash leading to traffic drop. PR1740390

Network Management and Monitoring

- EVO:JDI_EVO_REGRESSION:PLATFORM[ephemeral]:mgd core @ db_oneliner_perhaps. PR1753241
- On Junos OS Evolved platforms, SNMP cold start trap is observed on console log upon system reboot, but it is not sent out to external server. PR1788308

Routing Policy and Firewall Filters

- On all Junos OS Evolved PTX platforms, the fwstatsd process will crash when openconfig-NI filter last term is only routing-instance action with no IPv4/IPv6 match. The process crash will impact show commands and openconfig state sensor support. PR1788695

Routing Protocols

- ISIS TLV holds a maximum of 255 bytes of information per TLV. TLV contains data that has more than 255 bytes, current RFCs and drafts don't specify how to deal with multiple TLVs of same prefix. Hence it is advisable to associate bier subdomain to different lo0 TLV size exceed more than 255 bytes. It is not possible to define how many Bier subdomain will fit for given lo0 since it depend on other applicable subtlv associated with prefix. It is recommended to have advance planning/ calculation before putting multiple bier subdomains under a single lo0 address. IS-IS TLV size exceeds 255 bytes, router behavior is undefine and may lead to multiple network issues. PR1776786

User Interface and Configuration

- If its a scaled configuration with some configuration groups containing both wildcard and non-wildcard configuration, then the time it takes to commit the configuration is high. If upgrade is performed with this scaled configuration, then initial commit after upgrade might timeout and lead to a corrupt databases which could lead to mustd crash. To avoid this problem, move the wildcard configuration to another group.PR1804515
- Netconf: File copy with password less authentication is failing with RPC request. PR1769911

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.

General Routing

- Junos OS Evolved: A high rate of SSH connections causes a Denial of Service (CVE-2024-39562). PR1741624
- The legacy inet6 address seen under vmb0 while modifying mgmt-0 IPv6 address with dadfailed. PR1796934
- Interfaces take a long time to come up after reboot when configured in scaled IFL environment. PR1809423
- [timing] [ptp] PTX10008 PTP-PTP & PTP-1pps : Few frequencies shall be failing to meet the mask. PR1624478
- DNS resolution over a routing-instance fails. PR1733616
- PTX10001-36mr: no debug logs created post boot of the DUT. PR1746103
- FPC unreachable due to running out of Guid space. PR1756452
- All FPC Mezz board status shows "unknown" on Junos OS Evolved PTX10004, PTX10008, and PTX10016. PR1758265
- Interface queue statistics are not displayed on `show interfaces queue` CLI command. PR1760134
- The license-check can get restarted. PR1760259
- Unknown sensors are added in PFE on all Junos OS Evolved platforms. PR1765107
- Traffic blackholed due to hardware errors like FPC/SIB power fail or fabric Link errors. PR1766674
- The xinted generates syslog messages "service ssh, accept: Invalid argument (errno = 22)" with high CPU usage. PR1767072
- FPCs experiences crash and restart whenever the network encounters either an MPLS LSP flap or a LAG flap. PR1767747
- All Junos OS Evolved platforms - ifmon process 100% utilization. PR1768113
- PCS errors on Ethernet interface on certain PTX platforms running Junos OS Evolved. PR1768453
- FPC offline causes PTX10003-160C to reboot. PR1768610

- L2TPV3 load-balancing not working properly and create out of order packet flow. PR1769545
- The **evo-aftmand-bt[15138]: [t:15257] [Error] Jexpr: Invalid pfeId** error logs seen on all Junos OS Evolved platforms. PR1770432
- MPLS traffic flow might not be as expected after PFE restart. PR1770859
- The hwdre application restart will lead to non-functioning of GNMIC and memory components. PR1771597
- The `show interfaces extensive | no-more` command is taking a longer time to display the output. PR1773428
- [Junos OS Evolved] PTX10001-36MR/PTX10004/8/16 BT - Tail-dropped packets count is getting incremented on strict-high queue which is not expected. PR1773709
- The Control Board and FPC will restart if the optics present in the first three ports of PTX10001-36MR draw 50W or more power. PR1775320
- Duplicate IPv4 address detection error not logged on syslog for Junos OS Evolved platforms. PR1775981
- Interface stay in link DOWN state when using third party optics. PR1776596
- The orchestratord core dump during JSU. PR1776669
- The evo-aftmand-bt process crash is observed during an RE switchover. PR1776828
- Functionality provided by arp/ndp publish argument does not work on all Junos OS Evolved platforms. PR1776871
- PTX10001-36MR: epp_epc_intr_shmem_err seen in logs. PR1777003
- PTX10004, PTX10008, and PTX100016 Junos OS Evolved - after GRES, the backup RE BITS left over alarms is still in CM alarms. PR1777209
- LAG interfaces will take longer than usual to come up in a scaled scenario with ALB PR1777759
- [EVO] Log message for DDoS violation information shows default time and date wrong when its violation state is cleared by "clear ddos-protection protocols states". PR1778668
- The SRV6 traceroute with more than 3 SIDs does not work. PR1778946
- [Junos OS Evolved] Committed configuration files are not preserved post software version rollback operation. PR1779593
- CLI configuration statement to enable or disable auto recovery of l2-learning module in EVO platforms PR1779797

- Ungraceful FPC removal from chassis causing "hwdre" crash and process stops leading to FPC in fault state. PR1781493
- The rewrite-rule configured for the forwarding-class are not working as expected. PR1782536
- License key is not installed after USB upgrade, through `set system license keys key`. PR1783509
- Higher AE traffic convergence observed in ALB configured AE interface. PR1784498
- Performing back to back SIB offline results in context deadline exceeded error on PTX10016. PR1784766
- IFBD lookup Failure in DLU, packets needs to be dropped rather than learned through control IFL in packetio. PR1785084
- Difference in TOD between EEC and PTP FPGA. PR1787869
- PTX10000 Junos OS Evolved platforms with high multicast route changes can trigger multicast traffic queue drops. PR1789679
- The pfestatsd process might fail to restart when running out of file descriptors. PR1790095
- Junos OS Evolved : Fails to display SNMP object values on channelized interface. PR1790394
- Port-mirroring issue observed when adding/deleting interface with port-mirror configuration. PR1796517
- Interface input drops and PFE statistics info cell drops shows an incorrect large value with any configuration leading to IFD bounce. PR1796895
- PTX10001-36MR enters booting loop when a system reboot or software upgrade initiated reboot is performed. PR1799275
- MPLS payload traffic coming over EVPN-MPLS tunnel is dropped on PTX Junos OS Evolved platforms. PR1799760
- PTX Junos OS DDoS stats values of "arrival rate" and "max arrival rate" on System-wide and FPC showing larger values than expected. PR1801290
- Minor Host 0 Voltage Threshold Crossed reported on PTX10001 systems. PR1801330
- [Junos OS Evolved] "Host 0 Disk 1 Labelled incorrectly" alarm is sometimes set and cleared in 5 seconds. PR1801436
- IRB interface output filters will not work. PR1801716
- LACP goes down after adding native-vlan-id on AE with MACsec enabled on child links. PR1802071
- FFT shared bandwidth policer update not working (Day-1 issue). PR1812144

- Observed **MCNHMBB index availability** alarm when system is subjected to heavy PIM join/prune churn. PR1792740

Class of Service (CoS)

- Deletion of classifier/rewrite with import statement, along with some extra rules leads to the cosd process crash. PR1787101

EVPN

- In the EVPN-MPLS scenario traffic between CE devices will drop. PR1786959

Infrastructure

- Tunnel interface configuration crossing routing instances can cause fibd to abort. PR1788995

Interfaces and Chassis

- Traffic loss will be observed when an invalid IPv6 link-local address is configured. PR1774767
- On PTX10003 4x10GE/4x25GE interface drop the traffic in working lanes when new lane is configured with 400GE as neighbor interface. PR1810718

Network Management and Monitoring

- Junos OS Evolved, observing periodic error message "CMDOUT (error: error renaming temp state file /var/lib/logrotate.status.tmp)." PR1747722
- The lfmd fails to send notification about CRC error is seen on link. PR1769373
- CFM OAM status is **MEP status: Platform Unsupported** state in PTX10004 with LC1201 linecard. PR1777354

- The mib2d crash is observed on Junos OS Evolved platforms with duplicate SNMP request. PR1815524

Routing Policy and Firewall Filters

- [PTX10004 EVO] Maximum configurable value of policer if-exceeding "bandwidth-limit" is 100Gbps. PR1798975
- [Junos OS Evolved] - Firewall policer limits are adjusted on PTX10001-36MR to match the limits in other PTX10000 platforms. PR1800423
- [Junos OS Evolved] - Firewall policer limits are adjusted on PTX10003 to match the limits in other PTX10000 platforms. PR1804725

Junos OS Evolved Release Notes for QFX Series

IN THIS SECTION

These release notes accompany Junos OS Evolved Release 23.4R2 for QFX5130-32CD, QFX5130-48C, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240, QFX5700, QFX5700E-FEB switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New in 23.4R2-S1

IN THIS SECTION

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

Hardware

- **New QFX5130-48CM switch (QFX Series)**—In Junos OS Evolved Release 23.4R2-S1, we introduce the new Juniper Networks® QFX5130-48CM Switch. We also support this release on the QFX5130-48C switch. QFX5130-48CM is our first 1-U fixed form factor switch that is completely optimized for 100-Gigabit Ethernet (GbE) server connections. The QFX5130-48CM switch offers high-density 100GbE access ports in an SFP-DD form factor optimized for servers, along with high-density 400GbE ports in a QSFP-56 form factor optimized for easy uplinks to data centers. The QFX5130-48CM switch provides a throughput of 8 terabit per second (Tbps) by means of:

  - Forty-eight high-density 100GbE access ports that support SFP-DD transceivers optimized for servers.
  - Eight high-density 400GbE ports that support QSFP56 transceivers optimized for easy uplinks to the spine layer in data centers.
  - Support for Media Access Control Security (MACsec) feature.

  The QFX5130-48CM runs Junos OS Evolved. We've designed it to meet the needs of demanding data center environments such as high-performance computing and research networks and cloud and service provider data centers.

Table 2: QFX5130-48CM and QFX5130-48C Feature Support

Feature	Description
CoS	• CoS support on EVPN VXLANs. [See CoS Support on EVPN VXLANs.] • Support for priority-based flow control (PFC) of untagged traffic at Layer 3 using Differentiated Services Code Points (DSCP). [See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.]
Ethernet switching and bridging	• Support for Q-in-Q tunneling with a Service Provider Style configuration. [See Configuring Q-in-Q Tunneling.] • LLDP support. [See Device Discovery Using LLDP.] • Support for MAC move limit with EVPN-VXLAN. [See Understanding MAC Move Limiting.]
Forwarding options	• Support for port mirroring in EVPN-VXLAN environments. [See How to Configure Remote Port Mirroring for EVPN-VXLAN Fabrics.]
High availability	• Support for VRRP on Packet Forwarding Engine. [See VRRP Overview].

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
Interfaces	• Support for BGP flow specification (flowspec). [See BGP.] • Support for 48 SFP-DD and 8 QSFP-DD ports. Each switch also supports two 10-Gbps SFP+ ports. We support the following port configurations on each switch: • 48x100-Gbps / 50-Gbps / 25-Gbps / 10-Gbps on SFP-DD ports • 8x400-Gbps / 200-Gbps / 100-Gbps / 40-Gbps on QSFP-DD ports • 2x10-Gbps on SFP+ ports [See Port Settings (Interface Guide for Switches).] • Support for MACsec on physical interfaces —This platform supports MACsec in dynamic CAK mode with GCM-AES-128, GCM-AES-256, GCM-AES-XPN-128, and GCM-AES-XPN-256 encryption. This feature is supported on physical interfaces for switch-to-host and switch-to-switch links. [See Configuring MACsec.]
Junos telemetry interface (JTI)	• JTI streaming support for hardware Routing Engine-based sensors. Subscribe to `/components/sensor` to stream hardware operational stages. Statistics include Routing Engine, power supply unit (PSU), Control Board, Flexible PIC Concentrator (FPC), and PIC states. [See Junos YANG Data Model Explorer.]

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
Multicast	• Multicast Listener Discovery (MLD) snooping and integrated routing and bridging (IRB) stitching support. [See Understanding MLD Snooping.] • Support for multicast forwarding. [See Multicast Overview.] • IGMP snooping support. [See IGMP Snooping Overview .] • IGMP and MLD multicast snooping and IRB elaboration with make-before-break (MBB). [See IGMP Snooping Overview.] [See Understanding MLD Snooping.]

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
Network management and monitoring	• Support for sFlow. [See Overview of sFlow Technology.] • Support for port mirroring and analyzers. [See Understanding Port Mirroring and Analyzers.] • Inband Flow Analyzer (IFA) 2.0 transit node support. [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.] • IPsec support for OSPFv2 and OSPFv3. [See Overview of IPsec.] [See Configuring OSPF Authentication.] [See Configuring IPsec Security Associations.] • DHCP stateless relay MIB support. [See Enterprise-Specific MIBs for Junos OS Evolved.]
Protection against DDoS attacks	• Support for distributed denial of service (DDoS) protection, which is enabled by default. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.] and protocols (DDoS) (ACX Series, PTX Series, and QFX Series).]

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
Platform and infrastructure	• Platform resiliency support for hardware components of each FRU. If a failure is detected on a hardware component, Junos OS Evolved: • Logs the message to give clear indication of failure details, including time stamp, module name, and component name. • Raises or clears alarms if applicable. • Performs local actions, such as self-healing and taking the component out of service.
Precision Time Protocol	• Transparent clock support: • With or without VLAN encapsulation • With PTP over IPv4 • With PTP unicast or multicast • On LAG and (multichassis link aggregation) MC-LAG • On all physical, IRB, and aggregated Ethernet interfaces
Routing options	• Support for Unified Forwarding Table (UFT). [See Understanding the Unified Forwarding Table.]

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
Routing protocols	• Support for redistribution of IPv4 routes with IPv6 next hop into BGP. [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.] • Support to collect ON_CHANGE BGP routing information base (RIB) telemetry statistics and BGP neighbor telemetry with sharding. [See Telemetry Sensor Explorer.] • Support for maximum reference bandwidth increased to 4 Tbps for IGP protocols. [See reference-bandwidth (Protocols IS-IS).] [See reference-bandwidth (Protocols OSPF).] • Support to check for autonomous system (AS) matches in BGP policy AS paths without regular expressions. [See Improve the Performance of AS Path Lookup in BGP Policy.] • Support to strip or replace BGP private AS. [See Autonomous Systems for BGP Sessions.] • BGP Monitoring Protocol (BMP) local RIB monitoring support for all RIBs with sharding. [See BGP Monitoring Protocol]. [See loc-rib.] [See rib-list.] • Support for bootstrapping route-validation database from a local file. [See validation (Origin Validation for BGP).]

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
Routing policy and firewall filters	• Sharding support for conditional route manager. [See Routing Policy Match Conditions]. [See rib-sharding.] [See show policy conditions.] • Support for fast lookup of origin and neighbor ASs. [See policy-options.] [See policy-statement.] • Firewall filter support on Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions.] • Support for profiles to improve the firewall filter scale. [See Planning the Number of Firewall Filters to Create.] • EVPN-VXLAN firewall filtering and policing. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).]

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
System management	• Secure boot and secure BIOS support. [See Secure Boot.] • CLI-based hash and ECMP resilient hashing support. [See enhanced-hash-key.] [See ecmp-resilient-hash]. • Support for dynamic load balancing (DLB). [See enhanced-hash-key.] • Support to configure firewall filters and interfaces programmatically using JET APIs. [See Overview of JET APIs.]
Software installation and upgrade	• Zero-touch provisioning (ZTP) support. [See Zero Touch Provisioning.]
Services applications	• Support for DHCPv4 and DHCPv6 stateless relay. [See DHCP Relay Agent.]
Support for optics supported on QSFP-DD ports, SFP+ ports, SFP56-DD optics, and DAC cables	• To view the hardware compatibility matrix for optical interfaces, transceivers, and direct attach copper (DAC) cables supported , see the Hardware Compatibility Tool.

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
VPNs	• Support for EVPN Type 5 routes. [See Understanding EVPN Pure Type-5 Routes.] • Support for assisted replication (AR) integrated with optimized intersubnet multicast (OISM) in an EVPN-VXLAN edge-routed bridging (ERB) fabric. [See Assisted Replication Multicast Optimization in EVPN Networks.] [See Optimized Inter-Subnet Multicast in EVPN Networks.] • EVPN-VXLAN support with MAC-VRF routing instances. [See EVPN User Guide.] • Support for EVPN-VXLAN fabric with an IPv6 underlay. [See EVPN-VXLAN with an IPv6 Underlay.] [See Example: Configure an IPv6 Underlay for Layer 2 VXLAN Gateway Leaf Devices.] • Support for symmetric IRB with EVPN Type 2 routes. [See Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics.] [See irb-symmetric-routing.] • Support for MLDv1, MLDv2, and MLD snooping with OISM and AR in EVPN-VXLAN fabrics. [See Optimized Intersubnet Multicast in EVPN Networks.]

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Support to determine IRB interface state changes based on local and remote connectivity states in EVPN fabrics. [See Determine IRB Interface State Changes from Local and Remote Connectivity States in EVPN Fabrics.] [See interface-state.] [See network-isolation.] • Overlay and customer edge IP (CE-IP) address ping utility and traceroute support for EVPN-VXLAN. [See Understanding Overlay ping and traceroute Packet Support.] • Support to block asymmetric EVPN Type 5 routes. [See EVPN Type 5 Route with VXLAN encapsulation for EVPN-VXLAN.] [See ip-prefix-routes.] • Support for DHCP relay in an EVPN-VXLAN. [See DHCP Relay Agent over EVPN-VXLAN.] • Support for coexistence of EVPN Type 2 and Type 5 routes. [See EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN.] • Support to interconnect an EVPN-VXLAN in a data center to an EVPN-VXLAN control plane in a WAN using a gateway model. [See Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network.]

Table 2: QFX5130-48CM and QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Support for OISM in an EVPN-VXLAN fabric. [See Optimized Inter-Subnet Multicast in EVPN Networks.] • Support for Service Provider Style interface configuration on EVPN-VXLAN Layer 3 gateways. [See Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network.] • Overlapping VLAN support in EVPN-VXLAN fabrics on ERB overlay leaf devices. [See Overlapping VLAN Support Using VLAN Translation in EVPN-VXLAN Networks.] [See vlan-rewrite.]

What's New in 23.4R2

IN THIS SECTION

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

Hardware

- **New QFX5130-48C switch (QFX Series)**—Starting in Junos OS Evolved Release 23.4R2, we introduce the Juniper Networks® QFX5130-48C Switch. QFX5130-48C is our first 1-U fixed form factor switch that is completely optimized for 100GbE server connections. The QFX5130-48C switch offers high-density 100GbE access ports in a SFP-DD form factor optimized for servers, along with high-density 400GbE ports in a QSFP-56 form factor optimized for easy uplinks to data centers. The QFX5130-48C provides a throughput of 8 Tbps by means of:

  - Forty-eight high-density 100GbE access ports that support SFP-DD transceivers optimized for servers.
  - Eight high-density 400GbE ports that support QSFP56 transceivers optimized for easy uplinks to the spine layer in data centers.

  The QFX5130-48C runs Junos OS Evolved. We've designed it to meet the needs of demanding data center environments such as high-performance computing and research networks and cloud and service provider data centers.

  Table 3: QFX5130-48C Feature Support

Feature	Description
Class of service	CoS support on EVPN-VXLAN. [See CoS Support on EVPN VXLANs.]

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
Ethernet switching and bridging	• Support for Q-in-Q tunneling with a service-provider-style configuration. [See Configuring Q-in-Q Tunneling.] • LLDP support. [See Device Discovery Using LLDP.] • Support for MAC move limit with EVPN-VXLAN. [See Understanding MAC Move Limiting.]
Forwarding options	• Support for port mirroring in EVPN-VXLAN environments. [See How to Configure Remote Port Mirroring for EVPN-VXLAN Fabrics.]
High availability	• VRRP support on Packet Forwarding Engine. [See VRRP Overview].
Interfaces	• Support for BGP flowspec. [See BGP.]
Junos Telemetry Interface (JTI)	• JTI streaming support for hardware Routing Engine-based sensors. Subscribe to the / components/sensor to stream hardware operational stages. Statistics include Routing Engine, power supply units (PSUs), Control Boards, FPCs, and PICs states. [See Junos YANG Data Model Explorer.]

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
Multicast	• MLD snooping and IRB stitching support . [See Understanding MLD Snooping.] • Support for multicast forwarding. [See Multicast Overview.] • IGMP snooping support. [See PIM Overview.] • IGMP, MLD multicast snooping, and IRB elaboration with MBB. [See IGMP Snooping Overview.] [See Understanding MLD Snooping.]
Network management and monitoring	• Support for sFlow. [See Overview of sFlow Technology.] • Support for analyzers and port mirroring. [See Understanding Port Mirroring and Analyzers.]

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Inband Flow Analyzer (IFA) 2.0 transit node support. [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.] • IPsec support for OSPFv2 and OSPFv3. [See Overview of IPsec.] [See Configuring OSPF Authentication.] [See Configuring IPsec Security Associations.] • DHCP stateless relay MIB support. [See Enterprise-Specific MIBs for Junos OS Evolved.]
Protection against DDoS attacks	• Supports DDoS protection, which is enabled by default. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.] and protocols (DDoS) (ACX Series, PTX Series, and QFX Series).]
Platform and infrastructure	• Platform resiliency support for hardware components of each FRU. If a failure is detected on a hardware component, Junos OS Evolved: • Logs the message to give clear indication of failure details, including time stamp, module name, component name & failure details. • Raises or clears alarms if applicable. • Performs local action, such as self-healing and taking the component out of service.

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
Precision Time Protocol	• Transparent Clock support • With or without VLAN encapsulation • With PTP over IPv4 • With PTP unicast or multicast • On LAG and MC-LAG • On all physical, IRB, and AE interfaces
Routing options	• Support for Unified Forwarding Table (UFT). [See Understanding the Unified Forwarding Table.]
Routing protocols	• Support for redistribution of IPv4 routes with IPv6 next hop into BGP. [See Understanding Redistribution of IPv4 Routes with IPv6 Next Hop into BGP.] • Support for collect ON_CHANGE BGP RIB telemetry statistics and BGP neighbor telemetry with sharding. [See Telemetry Sensor Explorer.] • Support for maximum reference bandwidth increased to 4 TB for IGP protocols. [See reference-bandwidth (Protocols IS-IS).] [See reference-bandwidth (Protocols OSPF).] • Support for check for AS matches in BGP policy AS paths without regular expressions. [See Improve the Performance of AS Path Lookup in BGP Policy.]

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Support for stripping or replacing BGP private AS. [See Autonomous Systems for BGP Sessions.] • BMP local RIB monitoring support for all RIBs with sharding. [See BGP Monitoring Protocol]. [See loc-rib.] [See rib-list.] • Support for bootstrapping route-validation database from a local file. [See validation (Origin Validation for BGP).]
Routing policy and firewall filters	• Sharding support for conditional route manager. [See Routing Policy Match Conditions]. [See rib-sharding.] [See show policy conditions.] • Support for fast lookup of origin and neighbor autonomous systems (ASs). [See policy-options.] [See policy-statement.]

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Firewall filter support on Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions.] • Support for profiles to improve the firewall filter scale. [See Planning the Number of Firewall Filters to Create.] • EVPN-VXLAN firewall filtering and policing. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).]
System management	• Secure boot and secure BIOS support. [See Secure Boot.] • CLI-based hash and ECMP resilient hashing support. [See enhanced-hash-key.] [See ecmp-resilient-hash].
	• Support for dynamic load balancing (DLB). [See enhanced-hash-key.] • Support for configuring firewall filters and interfaces programmatically using JET APIs. [See Overview of JET APIs.]
Software installation and upgrade	• ZTP support. [See Zero Touch Provisioning.]

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
Services applications	• Support for DHCPv4 and DHCPv6 stateless relay. [See DHCP Relay Agent.]
Support for optics supported on QSFP-DD ports, SFP+ ports, SFP56-DD optics, and DAC cables	• To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported , see the Hardware Compatibility Tool.

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
VPNs	• Support for EVPN Type 5 routes. [See Understanding EVPN Pure Type-5 Routes.] • Assisted replication (AR) integrated with optimized intersubnet multicast (OISM) in an EVPN-VXLAN edge-routed bridging (ERB) fabric support. [See Assisted Replication Multicast Optimization in EVPN Networks.] [See Optimized Inter-Subnet Multicast in EVPN Networks.] • EVPN-VXLAN support with MAC-VRF routing instances. [See EVPN User Guide.] • Support for EVPN-VXLAN fabric with an IPv6 underlay. [See EVPN-VXLAN with an IPv6 Underlay.] [See Example: Configure an IPv6 Underlay for Layer 2 VXLAN Gateway Leaf Devices.] • Support for symmetric IRB with EVPN Type 2 routes. [See Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics.] [See irb-symmetric-routing.] • Support for MLDv1, MLDv2, and MLD snooping with OISM and AR in EVPN-VXLAN fabrics. [See Optimized Intersubnet Multicast in EVPN Networks.]

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Support for determining IRB interface state changes based on local and remote connectivity states in EVPN fabrics. [See Determine IRB Interface State Changes from Local and Remote Connectivity States in EVPN Fabrics.] [See interface-state.] [See network-isolation.] • Overlay and CE-IP ping and traceroute support for EVPN-VXLAN. [See Understanding Overlay ping and traceroute Packet Support.]

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	• Support for blocking asymmetric EVPN Type 5 routes. [See EVPN Type 5 Route with VXLAN encapsulation for EVPN-VXLAN.] [See ip-prefix-routes.] • Support for DHCP relay in EVPN-VXLAN. [See DHCP Relay Agent over EVPN-VXLAN.] • Support for coexistence of EVPN Type 2 and Type 5 routes . [See EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN.] • Support for Interconnecting EVPN-VXLAN in a data center to an EVPN-VXLAN control plane in a WAN using a gateway model. [See Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network.] • Support for OISM in an EVPN-VXLAN fabric. [See Optimized Inter-Subnet Multicast in EVPN Networks.] • Support for service-provider-style interface configuration on EVPN-VXLAN Layer 3 gateways. [See Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network.] • Overlapping VLAN support in EVPN-VXLAN fabrics on edge-routed bridging (ERB) overlay leaf devices.

Table 3: QFX5130-48C Feature Support *(Continued)*

Feature	Description
	[See Overlapping VLAN Support Using VLAN Translation in EVPN-VXLAN Networks.] [See vlan-rewrite.]

- **New QFX5230-64CD switch (QFX Series)**—Starting in Junos OS Evolved Release 23.4R2, QFX5230-64CD offers high-density 400-Gigabit Ethernet (GbE) access ports in a QSFP-DD form factor optimized for high-end spine and super-spine layer of the IP fabric multitier architecture in a 2-RU fixed form factor. The QFX5230-64CD switch provides a unidirectional throughput of 25.6 terabit per second (Tbps) and offers 64 400GbE network ports and up to 128x200GbE, 256x100GbE, 64x40GbE, 256x25GbE, and 258x10GbE ports.

Table 4: QFX5230-64CD Feature Support

Feature	Description
Chassis	• Support for inbuilt Routing Engine, Control Board, power supply unit, fan trays, Flexible PIC Concentrators (FPCs), and PICs on the QFX5230-64CD switch. • We have extended the Junos environment monitoring (EM) policy to include optics temperature sensors for QFX5230-64CD switches. It includes the following features: • The optics EM policy incorporates periodically polled temperature readings of optical modules in the system to automatically manage the fan speed. • 100GbE and 400GbE optics firmware automatically triggers optics shutdown when the high-temperature threshold is breached. • EM policy is enabled by default on all 100GbE and 400GbE optics interfaces, except for loopback optics and direct attach copper (DAC) cables. You can use the `set chassis fpc fpc_slot pic pic_slot port port_no no-temperature-monitoring` command to explicitly disable the EM policy on specific WAN ports. Use the `show chassis environment` command to view the optics temperature. [See temperature-sensor.]

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
CoS	• Support for CoS features on Layer 2 and Layer 3 interfaces. Both IPv4 and IPv6 unicast routing are supported. Other supported CoS features include: • Classification and rewrite rules for Differentiated Services code point (DSCP) and IEEE-802.1p. • Port scheduling • Shared buffer • Priority-based Flow Control (PFC) based on IEEE-802.1p. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2). • Weighted Random Early Drop (WRED) and Explicit Congestion Notification (ECN) • Telemetry support for CoS queue statistics exported using the sensor `/junos/system/linecard/qmon-sw/`. [See Traffic Management User Guide (QFX Series Switches and EX4600 Switches).]

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
EVPN	• Support for firewall filtering and policing on the Ethernet VPN–Virtual Extensible LAN (EVPN-VXLAN) network. [See Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).] • Support for sFlow on EVPN-VXLAN network. [See Overview of sFlow Technology.] • Support for port mirroring and analyzers on EVPN-VXLAN network. [See Port Mirroring and Analyzers in an EVPN-VXLAN Environment.]
Forwarding and sampling	• Support for dynamic load balancing (DLB) and resilient hashing (RH) for equal-cost multipath (ECMP) routes. DLB and RH are not supported on Link Aggregation Group (LAG). [See Dynamic Load Balancing, Use of Resilient Hashing to Minimize Flow Remapping, and ecmp-resilient-hash.]
High availability	• Support for Fast Fast Boot (FFB)-enabled in-service software upgrade (ISSU). [See Understanding In-Service Software Upgrade (ISSU).]

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Interfaces	• QFX5230-64CD has 64 QSFP56-DD ports and 2 SFP+ ports. The QSFP56-DD ports support the following speeds: • 400GbE • 200GbE • 100GbE • 50GbE • 40GbE The QSFP-DD ports also support the following speeds (with breakout cables): • 50GbE • 25GbE • 10GbE The SFP+ ports support 10GbE. QFX5230-64CD supports 1x400GbE, 2x200GbE, 4x100GbE, 2x100GbE, 1x100GbE, 2x50GbE, 1x50GbE, 1x40GbE, 4x25GbE, and 4x10GbE channelizations. [See Port Settings.]

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Layer 2 features	• Support for Layer 2 unicast forwarding and VRRP. [See Understanding VRRP.] • Support for IGMP snooping. This includes: • IGMP snooping with IGMPv1, IGMPv2, and IGMPv3 • IGMP proxy • IGMP querier at Layer 2 • Any-source multicast (ASM) and source-specific multicast (SSM) modes • Virtual router (VRF-lite) IGMP snooping • IGMP snooping with integrated routing and bridging (IRB) [See IGMP Snooping Overview, Multicast Overview, and Integrated Routing and Bridging.]

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Layer 3 features	• Support for DHCP stateless relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6. [See DHCP Relay Agent.] • Support for Layer 3 unicast forwarding and generic routing encapsulation (GRE) tunneling. We support both IPv4 and IPv6 unicast routing . [See Generic Routing Encapsulation (GRE).] • Support for Layer 3 multicast forwarding includes: • PIM first hop router (FHR) Rendezvous point (RP) functionality • MSDP • Make-before-break (MBB) support for multicast receivers on existing Layer 3 aggregated Ethernet (ae*x*) or link aggregation group (LAG) interfaces. Support includes member addition, member deletion, link up, and link down events. • PIM source-specific multicast (SSM) • PIM sparse mode (SM) • PIM dense mode (DM) • L3 multicast forwarding on integrated routing and bridging (IRB) interfaces: • IPv4 and IPv6 multicast • IGMP v1/v2/v3

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
	• Multicast Listener Discovery (MLD) v1/v2 • Any-source multicast (ASM) and source-specific multicast (SSM) modes [See Multicast Routing Protocols and PIM Overview.]
Network management and monitoring	• Support for sFlow. [See Overview of sFlow Technology.] • Support for port mirroring and analyzers. The QFX5230-64CD switches can support a maximum of eight port mirroring sessions. [See Understanding Port Mirroring and Analyzers.]

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Platform and infrastructure	• Platform resiliency support for hardware components of each FRU in a QFX5230-64CD switch. If a failure is detected on a hardware component, Junos OS Evolved: • Logs the message to give clear indication of failure details, including time stamp, module name, and component name. • Raises and clears alarms, if applicable. • Raises SNMP trap. • Makes the LED glow to indicate FRU fault, if an LED is present. • Performs local actions such as self-healing or taking the component out of service. • Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See Overview of JET APIs.]
Protection against DDoS attacks	• Supports configuration and installation of policers at the Packet Forwarding Engine (PFE) level for defense from DDoS attacks. By default, DDoS protection is enabled for many protocols on the QFX5230-64CD switches. [See Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers, show ddos-protection statistics, and show ddos-protection version.]
Routing policy and firewall filters	• Firewall filter support on Layer 2 and Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions and Configuring Enhanced Egress Firewall Filters.]

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Services applications	• Support for generic routing encapsulation (GRE) features: • GRE tunnels over GigE, LAG, and VLAN • Tagged sub-interfaces • Payload protocol for IPv4 and IPv6 • Delivery protocol for IPv4 • Multicast over GRE tunnels • Tunnel statistics • VRF with GRE • Time-to-live (TTL) [See Generic Routing Encapsulation (GRE).]

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Software installation and upgrade	• Firmware upgrade support. The following commands are supported: • `request system firmware upgrade fpc slot 0 bcm-pfe` • `request system firmware upgrade fpc slot 0 dpll` • `request system firmware upgrade fpc slot 0 dpll-cfg` • `request system firmware upgrade fpc slot 0 opticscpld<0|1|2>` • `request system firmware upgrade psm slot <0|1>` • `request system firmware upgrade re bios` • `request system firmware upgrade re fancpld` • `request system firmware upgrade re fpga` • `request system firmware upgrade re i210` • `request system firmware upgrade re ssd <disk1|disk2>` • `request system firmware upgrade re xmcfpga` [See request system firmware upgrade.] • Support for secure BIOS and secure boot implementation based on the UEFI 2.4 standard. [See Secure Boot.] • Zero Touch Provisioning (ZTP) support for WAN interfaces and DHCPv6 options. [See Zero Touch Provisioning.]

Table 4: QFX5230-64CD Feature Support *(Continued)*

Feature	Description
Support for optics	• 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.

- **New QFX5240 switches (QFX Series)**—Starting in Junos Evolved OS Release 23.4R2, we introduce the 800-Gigabit Ethernet (GbE) data center switches, QFX5240-64OD and QFX5240-64QD. These switches offer 64 800GbE OSFP and QSFP-DD ports. Using breakout cables, you can configure 64 ports of 800GbE, 128 ports of 400GbE, and 256 ports of 100GbE for QSFP5240-OD and QSFP5240-QD.

  Table 5: QFX5240 Feature Support

Feature	Description
Chassis	• Support for inbuilt Routing Engine, Control Board, power supply units, fan trays, Flexible PIC Concentrators (FPCs), and PICs on QFX5240-64OD and QFX5240-64QD switches.

Table 5: QFX5240 Feature Support *(Continued)*

Feature	Description
CoS	• Support for CoS features on Layer 2 and Layer 3 interfaces. Supported CoS features include: • IPv4 and IPv6 unicast routing. • Classification and rewrite rules for Differentiated Services code point (DSCP) and IEEE-802.1p • Port scheduling • Shared buffer • Priority-based Flow Control (PFC) based on IEEE-802.1p. DSCP-based PFC is required to support remote direct memory access (RDMA) over converged Ethernet version 2 (RoCEv2). • Weighted Random Early Drop (WRED) and Explicit Congestion Notification (ECN) • Telemetry support for CoS queue statistics exported using the sensor `/junos/system/linecard/qmon-sw/`. [See Traffic Management User Guide (QFX Series Switches and EX4600 Switches).]
Forwarding and sampling	• Support for dynamic load balancing (DLB) (for port speeds over 50Gbps) and resilient hashing (RH) for ECMP routes. DLB and RH are not supported on Link Aggregation Group (LAG) or when an LAG is one of the egress ECMP members. [See Dynamic Load Balancing, Use of Resilient Hashing to Minimize Flow Remapping, and ecmp-resilient-hash.]

Table 5: QFX5240 Feature Support *(Continued)*

Feature	Description
Interfaces	• The QFX5240 switches have 64x800GbE OSFP ports on QFX5240-64OD and 64x800GbE QSFP-DD ports on QFX5240-64QD. The last two ports (64 and 65) are 2x10GbE SFP on both the QFX5240 variants. The QFX5240-64OD and QFX5240-64QD ports support the following speeds: • 1x800 Gbps • 2x400 Gbps • 4x200 Gbps • 8x100 Gbps **NOTE**: On the QFX5240 switches, the runts (under Input errors) and fragment frames (under MAC statistics) counters do not increment in the output of the **show interfaces extensive** command. These counters are not supported due to a hardware limitation. [See Port Settings.]

Table 5: QFX5240 Feature Support *(Continued)*

Feature	Description
Layer 2 features	• Support for Layer 2 unicast forwarding and VRRP. [See Understanding VRRP.] • Support for IGMP snooping includes: • IGMP snooping with IGMPv1, IGMPv2, and IGMPv3 • IGMP proxy • IGMP querier at Layer 2 • Any-source multicast (ASM) and source-specific multicast (SSM) modes • Virtual router (VRF-lite) IGMP snooping • IGMP snooping with integrated routing and bridging (IRB) [See IGMP Snooping Overview, Multicast Overview, and Integrated Routing and Bridging.]

Table 5: QFX5240 Feature Support *(Continued)*

Feature	Description
Layer 3 features	• Support for DHCP stateless relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6. [See DHCP Relay Agent.] • Support for Layer 3 unicast forwarding and generic routing encapsulation (GRE) tunneling. We support both IPv4 and IPv6 unicast routing. [See Generic Routing Encapsulation (GRE).] • Support for Layer 3 multicast forwarding includes: • PIM first hop router (FHR) Rendezvous point (RP) functionality • Multicast Source Discovery Protocol (MSDP) • Make-before-break (MBB) support for multicast receivers on existing Layer 3 aggregated Ethernet (ae*x*) or link aggregation group (LAG) interfaces.Support includes member addition, member deletion, link up, and link down events. • PIM source-specific multicast (SSM) • PIM sparse mode (SM) • PIM dense mode (DM) • L3 multicast forwarding on integrated routing and bridging (IRB) interfaces: • IPv4 and IPv6 multicast • IGMP v1/v2/v3 • Multicast Listener Discovery (MLD) v1/v2 • Any-source multicast (ASM) and source-specific multicast (SSM) modes [See Multicast Routing Protocols and PIM Overview.]

Table 5: QFX5240 Feature Support *(Continued)*

Feature	Description
Network management and monitoring	• Support for sFlow. [See Overview of sFlow Technology.] • Support for port mirroring and analyzers. The QFX5240-64OD and QFX5240-64QD switches can support a maximum of seven port mirroring sessions. [See Understanding Port Mirroring and Analyzers.]
Platform and infrastructure	• Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See Overview of JET APIs.] • Platform resiliency support on QFX5240-64OD and QFX5240-64QD switches for hardware components of each FRU. If a failure is detected on a hardware component, Junos OS Evolved: 1. Logs the message to give clear indication of failure details, including time stamp, module name, and component name. 2. Raises SNMP trap. 3. Makes the LED glow to indicate FRU fault. 4. Performs local actions such as self-healing or taking the component out of service. [See QFX5240 Switch Hardware Guide and Chassis-Level User Guide.]

Table 5: QFX5240 Feature Support *(Continued)*

Feature	Description
Protection against DDoS attacks	• Supports configuration and installation of policers at the Packet Forwarding Engine (PFE) level for defense from distributed denial of service (DDoS) attacks. By default, DDoS protection is enabled for many protocols on the QFX5240-64OD and QFX5240-64QD switches. [See Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers, show ddos-protection statistics, and show ddos-protection version.]
Routing policy and firewall filters	• Firewall filter support on Layer 2 and Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions and Configuring Enhanced Egress Firewall Filters.]
Services applications	• Support for generic routing encapsulation (GRE) features: • GRE tunnels over GigE, LAG, and VLAN • Tagged sub-interfaces • Payload protocol for IPv4 and IPv6 • Delivery protocol for IPv4 • Multicast over GRE tunnels • Tunnel statistics • VRF with GRE • Time-to-live (TTL) [See Generic Routing Encapsulation (GRE).]

Table 5: QFX5240 Feature Support *(Continued)*

Feature	Description
Software installation and upgrade	• Firmware upgrade support. The following commands are supported: • `request system firmware upgrade re bios` • `request system firmware upgrade re fpga` • `request system firmware upgrade re gfpga` (QFX5240-64OD FPGA) • `request system firmware upgrade re ssd` • `request system firmware upgrade fpc opticscpld<0/1/2>` • `request system firmware upgrade re fancpld` • `request system firmware upgrade re i210` • `request system firmware upgrade fpc bcm-pfe` [See request system firmware upgrade.] • Support for zero-touch provisioning (ZTP) over IPv4 and IPv6 on the management and WAN interfaces. [See Zero Touch Provisioning.] • Support for USB booting. **NOTE**: On QFX5240 switches, only UEFI boot media (UEFI USB, UEFI NVME, UEFI network, and so on) is supported. You must select USB (UEFI USB) manually from the BIOS menu or use the `request node reboot re0 usb` command to boot from USB.
Support for optics	• 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.

- **New QFX5130E-32CD and QFX5700E-FEB models (QFX Series)**—

Starting in Junos Evolved OS Release 23.4R2, we introduce the QFX5130E-32CD and QFX5700E-FEB derivatives that are based on our QFX5130 and QFX5700 products. Both these new derivatives are based on the Trident 4-X11E platform.

Chassis

- **Platform software resiliency support (QFX5130-48C)**—Starting in Junos OS Evolved Release 23.4R2, platform software resiliency support is provided for QFX5130-48C, that includes:
  - CPU
  - Field replaceable units (FRUs)
  - Memory
  - USB port
  - Management Ethernet ports
  - Field-programmable gate array (FPGA) board
  - Optics panel
  - Fan tray
  - Power supply module

  If a failure is detected on a hardware component, the Junos OS Evolved software:
  - Logs the message with failure details, including time stamp, module name, and component name.
  - Raises or clears alarms, if applicable.
  - Makes the LED glow to indicate FRU fault.
  - Performs local action, such as self-healing and taking the component out of service.

  [See QFX5130 Switch Hardware Guide and Chassis-Level User Guide.]

- **Platform Software support for Routing Engine RE-QFX5130-48C (QFX5130-48C)**— Starting in Junos OS Evolved Release 23.4R2, platform software support is provided for Routing Engine RE-QFX5130-48C on QFX5130-48C. [See Chassis-Level User Guide.]

- **Port speed support (QFX5130E-32CD)**—Starting in Junos OS Evolved Release 23.4R2, you can configure 10-Gbps, 25-Gbps, 40-Gbps, or 100-Gbps port speed on QFX5130E-32CD switches.

  [See speed.]

- **Chassis management for FEB QFX5700E-FEB (QFX5700)**— Starting in Junos OS Evolved Release 23.4R2, you can use the following commands to manage the new Forwarding Engine Board QFX5700E-FEB and view information about the FEB:

  - `request chassis feb slot` *`slot-number`* `(offline | online | restart)` to offline, online, or restart the specified FEB

  - `show chassis feb` and `show chassis environment feb` commands to display FEB status information

  [See request chassis feb, show chassis feb, show chassis environment]

- **Optics EM policy support (QFX5230-64CD)**—Starting in Junos OS Evolved Release 23.4R2, we have extended the Junos Environment Monitoring (EM) policy to include optics temperature sensors for QFX5230-64CD switches. It includes the following features:

  - The Optics EM policy incorporates periodically polled temperature readings of optical modules in the system to automatically manage the fan speed.

  - 100GbE and 400GbE optics firmware automatically triggers optics shutdown when the high-temperature threshold is breached.

  - EM policy is enabled by default on all 100GbE and 400GbE optics interfaces, except for loopback optics and direct attach copper (DAC) cables.

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

  [See temperature-sensor.]

Junos Telemetry Interface

- **Configure an IP source address and routing instance for legacy gRPC dial-out connections (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Junos OS Evolved Release 23.4R2 supports configuring a source IP address and routing instance for legacy Remote Procedure Call (gRPC) service dial-out connections. In prior releases supporting legacy gRPC dial-out, the outgoing interface IP address is used as the source address without an option to configure a source IP address. This feature supports FLEX deployments, providing the ability to send dial-out from the router's specified IP address or interface address (such as a `loopback0` address).

  Use the `routing-instance` statement at the `[edit services analytics export-profile` *`profile-name`*`]` hierarchy level and the `local-address` *`ipv4 or ipv6 address`* statement at the `[edit services analytics export-profile` *`profile-name`*`]` hierarchy level.

  [See Using gRPC Dial-Out for Secure Telemetry Collection, routing-instance, and local-address.]

- **Telemetry for IPv4 and IPv6 traffic statistics (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Starting from Junos OS Evolved Release 23.4R2, you can stream IPv4 and IPv6

transit statistics using the native resource path **/junos/system/linecard/interface/traffic** or the OpenConfig resource path **/interfaces/interface/**. You can export the following fields are exported:

- **if_in_ipv4pkts**
- **if_in_ipv4_1sec_pkts**
- **if_in_ipv4_bytes**
- **if_in_ipv4_1sec_octets**
- **if_out_ipv4pkts**
- **if_out_ipv4_1sec_pkts**
- **if_out_ipv4_bytes**
- **if_out_ipv4_1sec_octets**
- **if_in_ipv6pkts**
- **if_in_ipv6_1sec_pkts**
- **if_in_ipv6_bytes**
- **if_in_ipv6_1sec_octets**
- **if_out_ipv6pkts**
- **if_out_ipv6_1sec_pkts**
- **if_out_ipv6_bytes**
- **if_out_ipv6_1sec_octets**

To enable transit statistics for the physical port, you must configure route accounting. To enable IPv4 route accounting or IPv6 route accounting, include the `route-accounting` statement at the `[edit forwarding-options family` *`family-name`*`]` hierarchy level.

[For sensors, see Junos YANG Data Model Explorer. For route accounting, see route-accounting.]

- **Telemetry for IPv4 and IPv6 traffic statistics (QFX5130-48C)**—Starting from Junos OS Evolved Release 23.4R2, you can stream hardware Routing Engine-based sensors using Junos telemetry interface (JTI). Subscribe to the **/components/** sensor to stream hardware operational states for the Routing Engines, power supply units (PSUs), control boards, FPCs, and PICs.

  [For sensors, see Junos YANG Data Model Explorer.

- **Telemetry support for CoS ingress packet drop accounting (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Junos OS Evolved Release 23.4R2 supports streaming counters for packets that are dropped due to ingress port congestion. Use the native sensor **/junos/system/linecard/**

interface/traffic to stream counters for priority-based flow control (PFC), explicit congestion notification (ECN), and ingress drops. Use the native sensor **/junos/system/linecard/qmon-sw/** to stream the priority group (PG) buffer utilization. You can stream counters for PFC, ECN, and ingress drops by means of OpenConfig using the sensor **/interfaces/interface/**.

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

Precision Time Protocol (PTP)

- **Transparent clock support (QFX Series)**—Starting in Junos OS Evolved Release 23.4R2, QFX5130E-32CD switches support the Precision Time Protocol (PTP) transparent clock feature.

  [See PTP Transparent Clock.]

- **Enterprise profile support (QFX Series)**—Starting in Junos OS Evolved Release 23.4R2, QFX5130E-32CD devices support the Precision Time Protocol (PTP) enterprise profile feature.

  [See PTP Enterprise Profile.]

Platform and Infrastructure

- **NIST purge method for media sanitization (QFX5130-32CD, QFX5220, and QFX5230-64CD)**—Starting in Junos OS Evolved Release 23.4R2, we've extended support for NIST media sanitization for SATA hard disk drives to:

  - Cryptographic scramble and block erase priorities for the purge method
  - Enhanced secure erase priority for the clear method

  For example, you can use this high level of data destruction when you pull a device from production. To maintain data security, you want to sanitize any disk drives in the device before it leaves your premises. The *NIST Special Publication 800-88* specifies the priority levels for sanitizing disk drives. In Junos OS Evolved, you sanitize a disk drive using the `request system zeroize (disk1 | disk2)` command. The sanitization process starts at the highest NIST sanitization priority that the disk drive supports. If that attempt fails, the process uses the method associated with the next lowest NIST priority level, and so on, until the disk is sanitized either using one of the NIST methods or using the Linux `dd` command.

  [See NIST Special Publication 800-88, Guidelines for Media Sanitization.]

Routing Protocols

- **BGP link bandwidth community (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5700, QFX5700E, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD)**—Starting in Junos OS Evolved Release 23.4R2, BGP can communicate link speeds to remote peers, enabling better optimization of traffic distribution for load balancing. A BGP group can send the *link-bandwidth* non-transitive extended community over an EBGP session for originated or received and readvertised link-bandwidth extended communities.

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

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

[See auto-sense, and group (Protocols BGP).]

- **Minimum ECMP (QFX5130-32CD, QFX5130-48C, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Starting in Junos OS Evolved Release 23.4R2, we support conditional advertising and withdrawal of BGP routes based on certain constraints such as bandwidth and minimum available next-hop ECMP. When a BGP receiver learns the same route from multiple BGP peers, BGP updates the active BGP path and the routing information base (RIB), also known as the routing table. The BGP export policy determines whether to advertise the BGP route to these next hops based on the number of ECMP BGP peers it receives the prefix from. A BGP route that has multiple ECMP BGP peers creates better resiliency in case of link failures. You can configure a BGP export policy to withdraw a BGP route unless it receives the BGP route prefix from a minimum number of ECMP BGP peers.

Services Applications

- **Protection against DDoS attacks (QFX5130-48C)**—Starting in Junos OS Evolved Release 23.4R2, you can configure and install policers at the Packet Forwarding Engine level for defense from DDoS attacks. DDoS protection is enabled by default on the QFX5130-48C switches.

  [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview and protocols (DDoS) (ACX Series, PTX Series, and QFX Series).]

Software Installation and Upgrade

- **Firmware upgrade support (QFX5130-48C)**—Starting in Junos OS Evolved Release 23.4R2, the QFX5130-48C switch supports the following firmware-upgrade commands:

  - `request system firmware upgrade fpc slot 0 bcm-pfe`
  - `request system firmware upgrade fpc slot 0 dpll`
  - `request system firmware upgrade fpc slot 0 dpll-cfg`
  - `request system firmware upgrade fpc slot 0 opticscpld<0|1|2|3>`
  - `request system firmware upgrade psm slot <0|1>`
  - `request system firmware upgrade re bios`

- `request system firmware upgrade re fancpld`
- `request system firmware upgrade re fpga`
- `request system firmware upgrade re i210`
- `request system firmware upgrade re ssd <disk1|disk2>`
- `request system firmware upgrade re xmcfpga`

[See request system firmware upgrade.]

- **Optimize reboot times by disabling default initialization and startup of certain L2 applications (QFX5130-32CD, QFX5130-48C, QFX5220, and QFX5700)**—Starting in Junos OS Evolved Release 23.4R2, when rebooting the device, the Layer 2 (L2) applications `l2ald`, `l2ald-agent`, `l2cpd`, and `l2cpd-agent` are initialized and started only if any of the following configuration hierarchy levels contain any configuration statements:
  - `[edit interface interface-name unit number family ethernet-switching]`
  - `[edit vlans]`
  - `[edit routing-instance instance-name instance-type virtual-switch]`
  - `[edit routing-instance instance-name instance-type mac-vrf]`
  - `[edit protocols l2-learning]`

  Additionally, `l2cpd`, and `l2cpd-agent` are initialized and started if the `[edit protocols lldp]` hierarchy level contains any configuration statements.

  As a result of this change, if your configuration already contains these configuration statements and you then delete all of them, these L2 applications stop running.

  [See request node reboot (re0 | re1) (Junos OS Evolved), request system reboot (Junos OS Evolved), and request system software add (Junos OS Evolved).]

Additional Features Optimized for AI-ML Fabrics

For more information about features optimized for AI-ML fabrics, see the AI-ML Data Center Feature Guide.

- **Reactive Path Rebalancing** (QFX5240)—Starting in 23.4R2 Junos OS Evolved Release, QFX5240 devices support Reactive Path Rebalancing. Reactive Path Rebalancing is an enhancement to the existing Flowlet mode in the Dynamic Load Balancing (DLB) feature. In the Flowlet mode of DLB, the user configures an inactivity interval. The traffic uses assigned outgoing interface until a pause in flow is greater than the inactivity timer. It is possible that the current outgoing link quality becomes worse over a period of time and the pause within the flow does not exceed the inactivity timer that

is configured. Classic Flowlet mode does not reassign to a different link within the inactivity interval and cannot utilize a better quality link. Reactive path rebalancing addresses this limitation by enabling the user to move the traffic to a link with a better quality in the Flowlet mode.

As per the existing DLB feature, each ECMP egress member link has a quality band assigned based on the traffic flowing through it. The quality band depends on the port load or number of egress bytes transmitted and queue buffer or the number of bytes waiting to be transmitted from the egress port. You can customize these attributes based on the traffic pattern flowing through the ECMP.

Benefits of the reactive path load balancing are:

- Optimal use of bandwidth
- Scalability
- Helps in avoiding load balancing inefficiencies due to long lived flows.

You need to configure DLB in the Flowlet mode. If you enable reactive path load balancing, packet reordering can occur when the flow moves from one port to another.

You need to satisfy the following rules to reassign a flow to a higher-quality member:

- An egress member port should be available whose quality is equal or greater than the current egress port.
- The packet random value is lower than the reassignment probability threshold value. When you configure a lower probability threshold value, flows move to higher-quality member at slower rate. For example, a probability threshold value of 200 has faster movement of macro flows to higher-quality member than probability threshold value of 50.

Example

Consider topology as shown in Figure 1 on page 131, where there are three ingress ports and two egress ports in a device. Also shown are table entries forwarding the traffic to each of the egress ports. All the ingress and egress ports are of the same speed.

Figure 1: Reactive Path Rebalancing

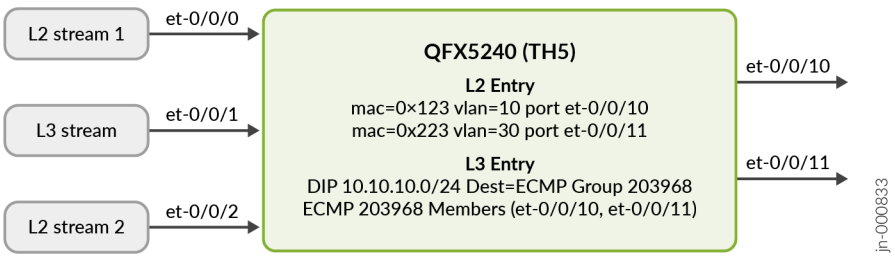

Reactive load rebalancing works with quality of delta 2 as follows:

1. Start stream 1 dmac 0x123 with rate 10 percent ingress port et-0/0/0 and egress out of et-0/0/10. Start stream 3 with rate of 50 percent ingress port et-0/0/1 and it egresses out of et-0/0/11.

   Egress link utilization is et-0/0/10 is 10 percent with Quality Band 6 and et-0/0/11 is 50 percent with quality band 5.

2. Start stream 2 dmac 0x223 with rate of 40 percent ingress port et-0/0/2 and it egresses out of et-0/0/11.

The reactive load balancing algorithm kicks in if the difference in quality bands for ports et-0/0/10 and et-0/0/11 is equal or higher than the configured delta of 2. The algorithm moves the stream 3 from et-0/0/11 to a better-quality member link, which is et-0/0/10 in this case.

After some time, you see et-0/0/10 link utilization of 60 percent with quality band of 5 as it egresses stream 1 and stream 3. The et-0/0/11 link utilization is of 40 percent with quality band of 5 as it egresses stream 2. See enhanced hash-key and show forwarding-options enhanced-hash-key.

- **Configurable FlowSet Table in DLB Flowlet Mode (QFX5000)**—DLB uses the FlowSet table to determine the egress interface of flows. The FlowSet table can hold a total of 32000 entries that needs to be distributed among 128 DLB ECMP groups. By default, these are divided equally allocating 256 entries per ECMP group. Starting in Junos 23.4R2 Junos OS Evolved Release, you can change the distribution of entries among the ECMP groups. If the FlowSet table has more ECMP group entries, the ECMP group can accommodate larger number of flows thereby achieving better flow distribution. See Dynamic Load Balancing.

- **PFC watchdog support (QFX5230-64CD, QFX5240-64OD, QFX5240-64QD)** —Starting in Junos OS Evolved Release 23.4R2, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD switches support the PFC watchdog feature. The PFC watchdog monitors PFC-enabled ports for PFC pause storms. When a PFC-enabled port receives PFC pause frames for an extended period of time and PFC watchdog does not detect flow control frames on that port, PFC watchdog mitigates the situation. It does this by disabling the queue where the PFC pause storm was detected for a configurable length of time called the recovery time. After the recovery time passes, PFC watchdog re-enables the affected queue.

  You configure PFC watchdog by including the `pfc-watchdog` statement at the `[class-of-service congestion-notification-profile` *`profile-name`*`]` hierarchy level. There are four parameters for PFC watchdog that you can configure for QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD switches:

  - `poll-interval`—The interval at which PFC watchdog checks the status of PFC queues, which can be 1, 10, or 100 milliseconds.

  - `detection`—The number of polling intervals the PFC watchdog waits before it mitigates the stalled traffic, from 1-15 intervals.

- `watchdog-action`—The action the PFC watchdog takes to mitigate a stalled traffic queue, either `drop` or `forward` all enqueued and newly arriving packets.
- `recovery`—How long the PFC watchdog disables the affected queue before it restores PFC on the queue, from 100-1500 milliseconds with a default of 200 milliseconds.

[See PFC Watchdog and congestion-notification-profile.]

- **Priority-based flow control X-ON Threshold support (QFX5230-64CD, QFX5240-64OD,and QFX5240-64QD)**—The priority-based flow control (PFC) X-ON threshold is the ingress port's priority group (PG) shared buffer limit. At this limit, the ingress port's peer resumes transmission of packets after a brief PAUSE because of the PFC message sent by this ingress port. You can fine tune the X-ON threshold through the congestion notification profile (CNP).

  [See xon (Input Congestion Notification).]

- **Per-queue alpha support (QFX5230-64CD, QFX5240-64OD,and QFX5240-64QD)**—You can tune globally the limit of buffers that each queue can consume from the shared pool based on the dynamic threshold setting called the alpha value. You can fine tune the alpha value on a per-queue basis through a scheduler.

  [See buffer-dynamic-threshold.]

- **Support for increased shared buffer pool (QFX5230-64CD, QFX5240-64OD,and QFX5240-64QD)**—By default, the QFX5230 switch allocates 73MB of the total 113MB of global buffer space to shared buffers, and the QFX5240 switch allocates 82MB of the total 165MB of global buffer space to shared buffers. These switches allocate the remaining buffer space to dedicated buffers (ingress and egress). You can decrease the global dedicated buffer space from the default value, effectively increasing the global shared buffer space to up to 106MB on the QFX5230 and 147MB on the QFX5240.

  You can also define a dedicated buffer profile to increase or decrease the dedicated buffer allocated to an individual port. This feature is particularly useful for decreasing dedicated buffer space on unused or down ports, thereby increasing dedicated buffer space available to active ports.

  [See Configuring Ingress and Egress Dedicated Buffers.]

- **egress-quantization** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD— Starting in Junos OS Evolved Release 23.4R2, you can modify port load and port queue metrics from their default values so that when dynamic load balancing is enabled, the metrics are used to determine an optimal link. Use the new `egress-quantization` CLI to configure the desired ratio of port load metric to port queue metric based on the traffic pattern.

  [See egress-quantization.]

- **`rdma-opcode` firewall filter match condition (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Starting in Junos OS Evolved Release 23.4R2, `rdma-opcode` and `rdma-opcode-except` firewall filter match conditions have been added to enable match on InfiniBand Base Transport header opcode.

  [See rdma-opcode.]

- **BGP Support for Global Load Balancing in DC Fabric (QFX5240)**—Starting in Junos OS Evolved Release 23.4R2, a route with multiple ECMP links is hashed onto several links for load balancing. In a DC fabric, hashing is unable to ensure even load distribution over all ECMP links, which might result in congestions on certain links and underutilization on other links. Dynamic load balancing helps to avoid congested links to mitigate local congestion. However, dynamic load balancing cannot address some congestions. For example, AI ML traffic that has elephant flows and lacks entrophy causes congestions in the fabric. In this case, global load balancing (GLB) helps to mitigate these congestions.

  In a CLOS network the congestions on the first two next hops impacts the load balancing decisions of the local node and the previous hop nodes triggering global load balancing. If the route has only one next-next-hop, a simple path quality profile is created. If the route has more than one next next-hop node then a simple path quality profile is created for each next next-hop node.

  To enable global load balancing, include the `global-load-balancing` statement at the `[edit protocols bgp]` hierarchy level. We have disabled this statement by default.

- **Extended sFlow Functionality Support** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—Starting in Junos OS Evolved Release 23.4R2, we've extended the sflow monitoring functionality to support the following features:

  - Export of sFlow sample packets via mgmt_junos interface.

    By default, the management Ethernet interface (usually named fxp0 or em0 for Junos OS, or re0:mgmt-* or re1:mgmt-* for Junos OS Evolved) provides the out-of-band management network for the device. Out-of-band management traffic is not clearly separated from in-band protocol control traffic. Instead, all traffic passes through the default routing instance and shares the default inet.0 routing table.

    Once you deploy the mgmt_junos VRF instance, management traffic no longer shares a routing table (that is, the default routing table) with other control traffic or protocol traffic in the system. Traffic in the mgmt_junos VRF instance uses private IPv4 and IPv6 routing tables.

    We've introduced a new configuration option "`routing-instance`" at `[edit protocol sflow collector]` hierarchy level to specify the routing instance name.

  - Export of sFlow sample packets via non-default VRF WAN ports.

  sFlow is a traffic monitoring protocol that supports VRFs. sFlow provides traffic sampling on configured ports based on sample rate and port information to a collector. An sFlow monitoring system consists of an sFlow agent embedded in the device and up to four external collectors. The

sFlow agent performs packet sampling and gathers interface statistics, and then combines the information into UDP datagrams that are sent to the sFlow collectors.

Collectors can be added and per VRF so that collectors can be spread out across different VRFs. The sFlow forwarding port can belong to a non-default VRF, and captured sFlow packets will have correct sample routing next hop information.

With this extended feature, an sFlow collector can be connected to the switch through the management network. The software forwarding infrastructure daemon (SFID) on the switch looks up the next-hop address for the specified collector IP address to determine whether the collector is reachable by way of the management network or data network.

Use the "`show sflow collector detail`" command to display the additional field "Routing Instance Name" to indicate the VRF name on which collector is reachable and "Routing Instance Id" that is corresponding to that VRF.

[See collector and show sflow collector.]

- **Per-queue accounting of explicit congestion notification (ECN) packets** (QFX5130, QFX5220, QFX5230, QFX5240, QFX5700)—Starting in Junos OS Evolved Release 23.4R2, counters on ECN-enabled queues increment when the queues experience congestion or receive packets that encountered congestion on another device. You can view these per-queue ECN accounting statistics through the `show interfaces queue` command. For example:

```
user@host# show interfaces queue et-0/0/5 forwarding-class network-control
Physical interface: et-0/0/5, up, Physical link is Up
...
    ECN-CE packets       :              8577686               482043 pps
    ECN-CE bytes         :           1252342156             70378315 bps
```

[See Understanding CoS Explicit Congestion Notification and show interfaces queue .]

- **SNMP support for PFC, ECN, and CoS ingress packet drop accounting** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—Junos OS Evolved Release 23.4R2 introduces SNMP support that helps to account for the packets that are dropped because of ingress port congestion. You can view and export the error counters data for explicit congestion notification (ECN), ingress drops, and priority-based flow control (PFC) using the following commands:

  - show snmp mib walk ifJnxTable
  - show snmp mib walk jnxCosPfcPriorityTable

  [See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved and show snmp mib.]

- **Telemetry support for CoS ingress packet drop accounting** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—Junos OS Evolved Release 23.4R2 supports streaming counters for packets that

are dropped due to ingress port congestion. Use the native sensor /junos/system/linecard/interface/ traffic to stream counters for priority flow control (PFC), explicit congestion notification (ECN), and ingress drops.

Use the native sensor /junos/system/linecard/qmon-sw/ to stream the priority group (PG) buffer utilization. You can also stream counters for PFC, ECN, and ingress drops by means of OpenConfig using the sensor /interfaces/interface/.

Counters for priority flow control (PFC), explicit congestion notification (ECN), and ingress drops are exported using the sensor /junos/system/linecard/interface/traffic.

Counters for PFC, ECN, and ingress drops are also exported using OpenConfig sensor /interfaces/interface/. Priority group (PG) buffer utilization is exported using the sensor /junos/system/linecard/qmon-sw/.

[See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]

- **Remote port mirroring to IPv4/IPv6 address (GRE encapsulation) with DSCP, source-address, and rate-limiting parameters (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Starting in Junos OS Evolved Release 23.4R2, you can configure DSCP, source-address, and rate-limiting parameters in your configuration for remote port mirroring to IPv4 or IPv6 addresses. You use remote port mirroring to copy packets entering a port or VLAN and send the copies to the IPv4 or IPv6 address of a device running an analyzer application on a remote network (sometimes referred to as “extended remote port mirroring”). The mirrored packets are GRE-encapsulated.

  You configure `source-address` or `source-ipv6-address`, `dscp`, and forwarding-class options—either in the analyzer configuration or the port-mirroring configuration—under these hierarchies, respectively:

  - `[edit forwarding-options analyzer instance` *`instance-name`* `output]`
  - `[edit forwarding-options port-mirroring instance` *`instance-name`* `family inet|inet6 output]`

  You configure the forwarding class and the shaping-rate option under the `class-of-service` hierarchy, as follows:

  - `set class-of-service forwarding-classes class` *`class-name`* `queue-num` *`queue-number`*
  - `set class-of-service interfaces` *`interface-name`* `scheduler-map` *`map-name`*
  - `set class-of-service scheduler-maps` *`map-name`* `forwarding-class` *`class-name`* `scheduler` *`scheduler-name`*
  - `set class-of-service schedulers` *`scheduler-name`* `shaping-rate` *`rate`*

  [See Port Mirroring and Analyzers.]

- **Strip and replace BGP private-AS path (QFX5230-64CD, QFX5240-64OD, QFX5240-64QD)**—Starting in Junos OS Evolved Release 23.4R2, we have introduced the `strip-as-path` policy option that removes the incoming autonomous system (AS) path as part of the import policy for a BGP session

and replaces the received autonomous system (AS) path with the receiving router's local AS number for the receiving session. Note that the local AS number may be different from the number configured under `autonomous system` in the `[edit routing-options]` hierarchy.

If you need to normalize externally injected routes, you can use this policy option for the incoming autonomous system (AS) path so that it may be used similarly to routes that originate solely within the fabric. The new `strip-as-path` policy option has no impact on the BGP export policy.

You can configure the `strip-as-path` option under `policy-options then` clause:

```
set policy-options policy-statement do-strip term a then strip-as-path
```

[See Autonomous Systems for BGP Sessions.]

Additional Features

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

- **EVPN VXLAN L2 GW, ARP suppression** (QFX5230-64CD)

  [See Understanding EVPN with VXLAN Data Plane Encapsulation and EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression.]

- **HTTP and TCP probe types for RPM** (QFX5130-32CD, QFX5130-48C, QFX5220, and QFX5230-64CD). You can now configure the `http-get`, `http-metadata-get`, and `tcp-ping` probe types for real-time performance monitoring (RPM) probes. You must configure the `offload-type none` statement to be able to commit the configuration.

  [See probe-server, probe-type, and rpm.]

- **Inband Flow Analyzer (IFA) 2.0 transit node support** (QFX5240-64OD and QFX5240-64QD)

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

- **Wake-On-Lan Targeted Broadcast support for EVPN-VXLAN** (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, and QFX5700)

  [See Targeted Broadcast and targeted-broadcast.]

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

What's Changed in 23.4R2-S5

IN THIS SECTION

Learn about changes in behavior and syntax in this release for QFX Series switches.

EVPN

- **DHCPv6 solicitation packets are flooded with no-dhcp-flood configured on an IRB (QFX5130 and QFX5700)**—In EVPN-VXLAN networks using ERB architecture and deployed as DHCP-RELAY's, IRB interfaces configured with `no-dhcp-flood` are flooding DCHPv6 solicitation packets. This issue is addressed in Junos OS Evolved Release 23.4R2-S5 and later.

  [See no-dhcp-flood.]

What's Changed in 23.4R2

IN THIS SECTION

Learn about changes in behavior and syntax in this release for QFX Series switches.

EVPN

- **Change in options and generated configuration for the EZ-LAG configuration IRB subnet-address statement**—With the EZ-LAG `subnet-address inet` or `subnet-address inet6` or

  at the `[edit services evpn evpn-vxlan irb` *`irb-instance`* hierarchy, you can now specify multiple IRB subnet addresses in a single statement using the list syntax `addr1 addr2 ?`. Also, in the generated configuration for IRB interfaces, the commit script now includes default `router-advertisement` statements at the `edit protocols hierarchy level` for that IRB interface.

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

- **Updates to syslog EVPN_DUPLICATE_MAC messages**—EVPN_DUPLICATE_MAC messages in the System log (syslog) now contain additional information to help identify the location of a duplicate MAC address in an EVPN network. These messages will include the following in addition to the duplicate MAC address:

  - The peer device, if the duplicate MAC address is from a remote VXLAN tunnel endpoint (VTEP).
  - The VLAN or virtual network identifier (VNI) value.
  - The source interface name for the corresponding local interface or multihoming Ethernet segment identifier (ESI).

  For example: `Feb 27 22:55:13 DEVICE_VTEP1_RE rpd` *`39839`*: EVPN_DUPLICATE_MAC: MAC address move detected for 00:01:02:03:04:03 within instance=evpn-vxlan on VNI=100 from 10.255.1.4 to ge-0/0/1.0.

  For more on supported syslog messages, see System Log Explorer.

- **Limit on number of IP address associations per MAC address per bridge domain in EVPN MAC-IP database**—By default, devices can associate a maximum of 200 IP addresses with a single MAC address per bridge domain. We provide a new CLI statement to customize this limit, `mac-ip-limit` statement at the **edit protocols evpn** hierarchy level. In most use cases, you don?t need to change the default limit. If you want to change the default limit, we recommend that you don?t set this limit to more than 300 IP addresses per MAC address per bridge domain. Otherwise, you might see very high CPU usage on the device, which can degrade system performance.

  [See mac-ip-limit.]

Infrastructure

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

  [See Path MTU Discovery.]

Junos OS API and Scripting

- **<get-trace> RPC support removed (ACX Series, PTX Series, and QFX Series)**—The `show trace application` *app-name* operational command and equivalent <get-trace> RPC both emit raw trace data. Because the <get-trace> RPC does not emit XML data, we've removed support for the <get-trace> RPC for XML clients.

- **XML output tags changed for `request-commit-server-pause` and `request-commit-server-start` (QFX Series)**—We've changed the XML output for the `request system commit server pause` command (`request-commit-server-pause` RPC) and the `request system commit server start` command (`request-commit-server-start` RPC). The root element is <commit-server-operation> instead of <commit-server-information>, and the <output> tag is renamed to <message>.

Network Management and Monitoring

- **NETCONF <copy-config> operations support a `file://` URI for copy to file operations (QFX Series)**—The NETCONF <copy-config> operation supports using a `file://` URI when <url> is the target and specifies the absolute path of a local file.

  [See <copy-config>.]

- **Device family identifier changed for native YANG modules (QFX Series)**—Starting in Junos OS Evolved Release 23.4R2, native YANG modules for QFX Series devices by default use the `junos` device family identifier instead of the `junos-qfx` identifier in the module's name, namespace, and filename. With this change, all devices running Junos OS Evolved use the `junos` device family identifier. To emit device-specific modules that use the `junos-qfx` device family identifier, configure the `device-specific` and `emit-family-ns-and-module-name` statements at the `[edit system services netconf yang-modules]` hierarchy level.

  [See Understanding Junos YANG Modules.]

Routing Protocols

- **Optimized mesh group routes (QFX5110, QFX5120, QFX5130, QFX5700 and ACX Series)**— `show route snooping` for inet.1/inet6.1 table and `show route snooping table inet.1/inet6.1` displays only CE mesh group routes for platforms that support EVPN-MPLS or EVPN-VxLAN multicast. In earlier releases, other mesh groups like the VE mesh group were also displayed.

System Management

- **Additional `Upgrade` fields for the `show system applications detail` command (ACX Series, PTX Series, and QFX Series)**—The `show system applications detail` command and corresponding RPC include additional `Upgrade` output fields. The fields provide information about notifications and actions related to various upgrade activities.

  [See show system applications (Junos OS Evolved).]

Software Installation and Upgrade

- **`configuration` and `no-configuration` options for the `request system snapshot` command (QFX Series)**—When you omit or include the `configuration` option, the `request system snapshot` command copies the **/config** directory and the configuration stored for each installed software version to the alternate solid-state drive (SSD) as part of the snapshot. You can use the `no-configuration` option to exclude the **/config** directory and the configuration stored for each installed software version from the snapshot.

VPNs

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

  [See min-rate.]

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

  [See min-rate.]

Known Limitations

IN THIS SECTION

Learn about limitations 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

- IPv4 packets dropped due to egress MTU mismatch are not shown in the `show interfaces extensive output` in the CLI . To get these dropped packet statistics, the debug commands from the forwarding process must be used. PR1768752
- On QFX5240-64OD/64QD platform, the error statistics on the management port re0:mgmt-0 does not reflect the errors experienced including link flaps, FCS erorrs etc. PR1791886
- On QFX5240 and QFX5230 platforms with Junos OS Evolved 23.4R2, when PFC watchdog feature configured with recovery action of **drop** and if PFC storm continuously detected and recovered, CRC error counter could increase with small number under `show interface extensive` *`interface`* output. These CRC errors could occur only for larger frames (mtu>1024) received at high rate. Packet drops are seen on the interface due to PFC watchdog action of drop, which is expected. CRC errors are not seen with PFC watchdog recovery action of **forward**. PR1807420
- On QFX5700E, **pci 0000:xx:xx.x BAR xx: failed to assign** messages are seen during boot in the logs. These messages have no impact to functionality. PR1807706
- Storm control does not work on multicast packets on QFX5130, QFX5700 platform if broadcast packets are excluded from the storm control profile. To support IGMP, MLD snooping functionality, we've added flood in vlan rules in ASIC pipeline which causes mcast packets to be flooded as broadcast packets due to BCM TD4 ASIC limitation. PR1813514

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

- On QFX5230-64CD platform, 400G DAC cable of 2.5m length and 4X100G DAC BO might not link up with some peer devices. This issue is not seen with all peer devices. If this happens, please replace the 2.5m cable with a 1m DAC cable or use supported 400G Optics instead. PR1747315
- On QFX5130-48C, when image upgrade is done using CLI, the log messages will have kernel trace reported. There is no functionality impact due to the messages.PR1755406
- QFX5230-64CD: **FPC 0 idtRC38612 PLL LOCK Failure** major alarm is seen after the system comes up post a device reboot. This alarm is cleared after a few minutes once the system reaches steady state. No user action is required.PR1766984
- Whenever a QFX5130-48C device is reset using rear panel reset button by pressing the button for short or long duration, the port LEDs on ports 56/57 might glow in amber for a few seconds during boot up. The port LEDs amber status turn off and reflect correct port state once device is up.PR1792619
- The recommended procedure for FPC card removal is to do a graceful OIR by following the documented steps. However, if the FPC card is ungracefully jacked out or plugged in several times in a short interval, the software might show the FPC status as "FAULT". Following are the steps to recover from the fault state: PR1799333

  1. Jack out the FPC.
  2. Wait for 1 minute for FPC state to become **Fault** in `show chassis fpc pic-status` command.

3. Jack In FPC.
4. Wait for 1 minute for FPC state shows as present in `show chassis fpc pic-status` command.
5. Run CLI command `request chassis fpc online slot` to bring FPC in online state.
6. Wait for 3 minutes for all interfaces to come up.

- QFX5130E-32CD supports class A performance with default FEC applicable for the particular interface.PR1800144
- On a QFX5700 chassis with QFX5700E-FEB, during FEB offline or online operation, a power fault message is reported on the console. Following is an example of the message printed on to the console. **pcieport 0000:41:09.0: pciehp: Slot(9): Power fault**. PR1802959
- When a VXLAN encapsulated IP packet, or an IP packet with UDP port matching the VXLAN UDP port, is received on a vlan-tagging enabled interface, the switch drops the frame. This issue is not seen if the incoming port is an untagged interface, or if the interface is actually doing VXLAN encap/decap operations. In such cases, the device forwards the frame correctly. PR1805922
- When the QFX5700 is power cycled, in some cases, during the reboot of the system, the evoaft-jvisiond-brcm process might not start automatically. The operating system stops the service and generates a major alarm. This does not impact other forwarding or chassis functionality. Restarting the process through the CLI login shell clears the alarm and the process functionality is restored. PR1807624
- IPv4/IPv6 reserved multicast and L2 multicast traffic received over VXLAN access port is flooded out of all ports of the VXLAN except vtep. PR1811158
- On the QFX5220, QFX5130, QFX5700, QFX5230, or QFX5240 platforms, after clearing queue counters for the aggregated Ethernet interface from CLI, if one of the aggregated Ethernet member interfaces is immediately removed from aggregated Ethernet bundle, it displays invalid queue counters for the aggregated Ethernet interface. Clearing the queue counters from CLI again helps to resolve the issue.PR1811575
- No support for Layer 3 vlan-tagged sub-interfaces on VXLAN NNI side in Junos OS Evolved 23.4R2 on QFX5230. Only layer 3 interfaces or IRB can be configured on VXLAN NNI side. PR1813069
- When a QFX5700E chassis is powered off, there might be CPU kernel trace messages printed onto the console during power off. These messages have no impact to functionality.PR1813256
- QFX130E-32CD: Sometimes system fails to boot up during power cycle when there is insufficient time gap between power-off and power-on. The recommendation is to give a 120 seconds gap between power-off and power-on. PR1813294
- QFX5230-64CD, QFX5130-48C, : Help string for MTU settings on management interface are not correct. The range is showed as 256-9408 instead of 256-9216.PR1813591

- In rare cases, **FPC 0 Volt Sensor Fail" or "FPC 0 Voltage Threshold Crossed** alarm is seen after bootup on QFX5130-48C or QFX5130-48CM. PR1816064
- QFX5700 MacSec: Minor packet drops (< 0.0000001%) observed when MacSec is enabled. PR1816407
- USB disks with Junos OS Evolved images from Junos OS Evolved 23.4R2 onwards might not be detectable by Windows. They still have valid images, and can be used for Junos OS Evolved installs. The only issue is that new images cannot be installed on these USB disks because Windows no longer recognizes these USB drives. PR1819846
- On QFX5220, QFX5230, QFX5240, QFX5130, or QFX5700 platforms, for lossless queues packet admission control is based on ingress priority group thresholds only. Once lossless packet is accepted by ingress, there is no any further admission control checks based on egress queue thresholds. So dynamic threshold (alpha) setting configured at egress queue level, does not make any impact for lossless queues. For lossless queue, peak shared buffer utilisation is always based on the ingress priority group dynamic threshold. Do not configure per queue alpha through **buffer-dynamic-threshold** for lossless queues. On QFX5220, QFX5230, QFX5240, QFX5130, or QFX5700 platforms, ECN threshold is static. However PFC thresholds are dynamic. For effective DCQCN behaviour, make sure ECN threshold is always lower than the PFC XOFF threshold. Specifically when multiple ingress and egress ports involved in lossless queue congestion, there is high chance for PFC XOFF threshold hitting before ECN threshold. Proper ECN marking thresholds can be arrived by monitoring the peak buffer utilisation of congested lossless queues and fine tuning the ECN thresholds accordingly. PR1820266
- QFX5130-48CM: MacSec: Packet drops are seen when MKA session switches from primary to fallback In case the primary pre-shared-key (PSK) fails to establish a connection, the fallback PSK is used. Traffic loss is observed during switch from primary to the fallback session. No traffic drop is observed if the primary PSK is kept active without switching to the fallback PSK. Thus, as long as the primary PSK is active, the MACSec sessions pass traffic properly without any drop. PR1820549

Interfaces and Chassis

- Unified ISSU from Junos OS Evolved 23.2R2 to newer releases generates core files. Object data structure is extended for a fix and during the extension, the extensions between the releases should match. However, there is a mismatch with respect to such extensions between the releases, which causes the issue. PR1803068

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

- The `show interfaces extensive | no-more` command takes a longer time to display the output. PR1773428
- On Junos OS Evolved platform, SNMP walk fails to display SNMP object values on channelized interface. PR1790394
- QFX5230-64CD, QFX5240-64OD, QFX5240-64QD: USB stick continues to show up under `show chassis hardware detail` even after unplugging from the device. PR1793934
- IPv6 traffic statistics shows large or invalid values at times when route-accounting is deactivated and activated. PR1798636
- Network reachability issues, that is traffic drops when MAC moves from VGA mac and Remote IRB MAC to access interface. PR1805685
- Due to clocking issue in the external PHY, traffic drops are observed having MACsec enabled. PR1806957

Interfaces and Chassis

- On Junos OS Evolved QFX5700 platforms, the vmcore might be seen if any component is forcefully removed from the PCI (Peripheral Component Interconnect) bus. PR1739142

- The `monitor interface` or `show interface` does not display the interface description. PR1762065

Upgrade Your Junos OS Evolved Software

Products impacted: ACX7024, ACX7024X, ACX7100-32C, ACX7100-48L, ACX7332, ACX7348, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, PTX10016, PTX10002-36QDD, QFX5130-32CD, QFX5130-48C, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240, QFX5700, and QFX5700E.

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.

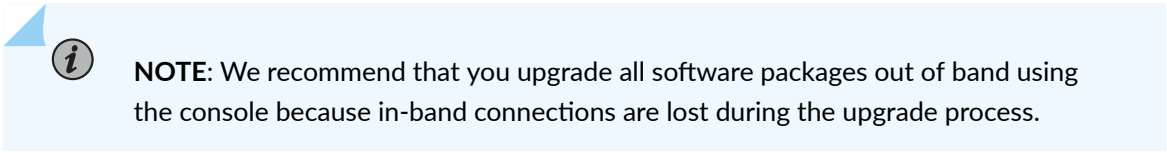

Upgrade and Downgrade Support Policy for Junos OS Evolved Releases

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

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

For both Standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 21.4 to the next three releases – 22.1, 22.2 and 22.3 or downgrade to the previous three releases – 21.3, 21.2 and 21.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 21.4 is an EEOL release. Hence, you can upgrade from 21.4 to the next two EEOL releases – 22.2 and 22.4 or downgrade to the previous two EEOL releases – 21.2 and 20.4.

Table 6: EOL and EEOL Releases

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

For more information about 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

  NOTE: To obtain information about the components that are supported on the devices and the special compatibility guidelines with the release, see the Hardware Guide for the product.

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

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

Requesting Technical Support

IN THIS SECTION

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/

- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/

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

Creating a Service Request with JTAC

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

- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

Revision History

23 September 2025—Revision 13, Junos OS Evolved Release 23.4R2

02 September 2025—Revision 12, Junos OS Evolved Release 23.4R2

17 July 2025—Revision 11, Junos OS Evolved Release 23.4R2-S5

11 November 2024—Revision 10, Junos OS Evolved Release 23.4R2-S1

14 October 2024—Revision 9, Junos OS Evolved Release 23.4R2

23 September 2024—Revision 8, Junos OS Evolved Release 23.4R2-S1

17 September 2024—Revision 7, Junos OS Evolved Release 23.4R2

29 August 2024—Revision 6, Junos OS Evolved Release 23.4R2-S1

14 August 2024—Revision 5, Junos OS Evolved Release 23.4R2

13 August 2024—Revision 4, Junos OS Evolved Release 23.4R2

31 July 2024—Revision 3, Junos OS Evolved Release 23.4R2

23 July 2024—Revision 2, Junos OS Evolved Release 23.4R2

10 July 2024—Revision 1, Junos OS Evolved Release 23.4R2