

Release Notes

Published
2025-12-23

Junos OS Evolved Release 25.2X100-D10

Introduction

Use these release notes to find new features, software limitations, and open issues for Junos OS Evolved Release 25.2X100-D10.

For more information on this release of Junos OS Evolved, see [Introducing Junos OS Evolved](#).



NOTE: Junos OS Evolved 25.2X100-D10 is a controlled release available only on the following platforms:

- QFX5130-32CD
- QFX5130E-32CD
- QFX5130-48C
- QFX5130-48CM
- QFX5220-32CD or 128C
- QFX5230-64CD
- QFX5240-64OD
- QFX5240-QD
- QFX5241-32OD
- QFX5241-64OD or QFX5241-64QD
- QFX5700
- QFX5700E

If you are looking for this release, contact your Juniper Networks Account Team for more information.

Table of Contents

Junos OS Evolved Release Notes for QFX Series

What's New | 1

Hardware | 2

Class of Service | 18

EVPN | 19

Forwarding Options | 19

Interfaces | 20

Junos Telemetry | 21

Layer 2 VPN | 22

Network Management and Monitoring | 22

Routing Protocols | 23

Additional Features | 24

What's Changed | 26

Known Limitations | 27

Open Issues | 27

Licensing | 27

Finding More Information | 28

Requesting Technical Support | 29

Revision History | 30

Junos OS Evolved Release Notes for QFX Series

IN THIS SECTION

- [What's New | 1](#)
- [What's Changed | 26](#)
- [Known Limitations | 27](#)
- [Open Issues | 27](#)

These release notes accompany Junos OS Evolved Release 25.2X100-D10 for QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220-32CD, QFX5220-128C, QFX5230-64CD, QFX5240-OD, QFX5240-QD, QFX5241-32OD, QFX5241-64OD, QFX5241-64QD, QFX5700, and QFX5700E switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

What's New

IN THIS SECTION

- [Hardware | 2](#)
- [Class of Service | 18](#)
- [EVPN | 19](#)
- [Forwarding Options | 19](#)
- [Interfaces | 20](#)
- [Junos Telemetry | 21](#)
- [Layer 2 VPN | 22](#)
- [Network Management and Monitoring | 22](#)
- [Routing Protocols | 23](#)
- [Additional Features | 24](#)

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

Hardware

- **QFX5241-32OD (QFX Series)**—The Juniper Networks® QFX5241-32OD Switches are fixed-configuration devices with 32 octal small form-factor pluggable (OSFP) ports that support speeds of up to 800 Gbps. Features such as 25.6 terabits per second (Tbps) throughput and 1-U shallow buffer design make these switches optimal as end-of-row, leaf, or spine devices in IP fabric architectures. The switches support 2400-watt (W) AC power supply units (PSUs) and front-to-back airflow.

To install the QFX5241-32OD switch and perform initial configuration, routine maintenance, and troubleshooting, see the [QFX5241-32OD Switch Hardware Guide](#). See [Feature Explorer](#) for the complete list of features for any platform.

Table 1: QFX5241-32OD Feature Support

Feature	Description
CoS	<ul style="list-style-type: none">• Support for CoS features on Layer 2 and Layer 3 interfaces, including:<ul style="list-style-type: none">• 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 for VLAN-tagged traffic. Protocols such as remote direct memory access (RDMA) over Converged Ethernet version 2 (RoCEv2) require DSCP-based PFC at Layer 3 for untagged traffic.• Weighted random early detection (WRED) and explicit congestion notification (ECN)• Telemetry support for CoS queue statistics exported using the sensor <code>/junos/system/linecard/qmon-sw/</code>. <p>[See Traffic Management User Guide (QFX Series Switches and EX4600 Switches).]</p>

Table 1: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
EVPN	<ul style="list-style-type: none"> Support for Layer 2 gateway and Address Resolution Protocol (ARP) suppression on EVPN-VXLAN. [See Understanding EVPN with VXLAN Data Plane Encapsulation, EVPN Proxy ARP and ARP Suppression, and overlay (Packet Forwarding Options).] Support for CoS, and firewall filtering and policing on EVPN-VXLAN. [See CoS Support on EVPN VXLANs and Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).] Support for Wake-on-LAN (WOL) targeted broadcast on EVPN-VXLAN. [See Targeted Broadcast and targeted-broadcast.] Support for EVPN-VXLAN Layer 2 gateway, including: <ul style="list-style-type: none"> Multihoming ARP suppression Layer 3 IPv4 underlay with integrated routing and bridging (IRB) and LAG Core isolation Broadcast, unknown unicast, and multicast (BUM) traffic forwarding by ingress replication only MAC move limits [See Understanding EVPN with VXLAN Data Plane Encapsulation, EVPN Proxy Arp and Arp Suppression, and Proxy NDP and NDP Suppression, IP Fabric Underlay Network Design and Implementation, overlay-ecmp, Edge-Routed Bridging Overlay Design and Implementation, Layer 2 Interface Status Tracking and Shutdown Actions for EVPN Core Isolation Conditions, and mac-move-limit.] Support for EVPN-VXLAN Layer 3 gateway, including:

Table 1: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> • Layer 3 VXLAN gateway in edge-routed bridging (ERB) fabric • Up to 256 VLANs with IRB enabled • Layer 3 underlay that supports IRB and LAG • ECMP in the underlay • IPv4 and IPv6 virtual gateway MAC address support for IRB interfaces • In-service software upgrade (ISSU) for Layer 3 gateway functionality <p>[See Understanding EVPN with VXLAN Data Plane Encapsulation, Example: Configuring an EVPN-VXLAN Edge-Routed Bridging Fabric with a Virtual Gateway, Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network, and IP Fabric Underlay Network Design and Implementation.]</p> <ul style="list-style-type: none"> • Support for EVPN-VXLAN Type 5 stitching, including: <ul style="list-style-type: none"> • Overlay and underlay ECMP • Type 5 stitching • Type 2 and Type 5 route coexistence • Symmetric IRB • ISSU <p>[See Understanding EVPN with VXLAN Data Plane Encapsulation, IP Fabric Underlay Network Design and Implementation, overlay-ecmp, Understanding EVPN Type 5 Routes, EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN, NSR and Unified ISSU Support for EVPN, and irb-symmetric-routing.]</p> <ul style="list-style-type: none"> • Support for sFlow technology on EVPN-VXLAN.

Table 1: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
	<p>[See Overview of sFlow Technology.]</p> <ul style="list-style-type: none">• Support for port mirroring and analyzers on EVPN-VXLAN. <p>[See Port Mirroring and Analyzers in an EVPN-VXLAN Environment.]</p> <ul style="list-style-type: none">• Support for forwarding EVPN data traffic on the spine device without any traffic loss while the leaf device performs a unified ISSU in spine-and-leaf topologies with external BGP (EBGP) connections. <p>[See Understanding Unified ISSU.]</p>

Table 1: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
Features optimized for AI-ML fabrics	<ul style="list-style-type: none"> Support for priority-based flow control (PFC) watchdog. [See PFC Watchdog and congestion-notification-profile.]
	<ul style="list-style-type: none"> Telemetry support for streaming IPv4 and IPv6 transit statistics using the native resource paths <code>/state/interfaces/interface[name='']/counters/ipv4/</code> and <code>/state/interfaces/interface[name='']/counters/ipv6/</code>. [See Junos YANG Data Model Explorer and route-accounting.]
	<ul style="list-style-type: none"> Support for enabling or disabling dynamic load balancing (DLB). You can use the <code>dynamic-load-balance</code> statement to selectively enable or disable DLB based on <code>rdma-opcode</code> match or any match available in firewall filters. The optimal link is determined based on the modified port load and port queue metrics when DLB is enabled. [See rdma-opcode, dynamic-load-balance-selective, and egress-quantization.]
	<ul style="list-style-type: none"> Support for PFC using Differentiated Services code points (DSCPs) at Layer 3 for untagged IPv6 traffic. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over Converged Ethernet version 2 (RoCEv2). [See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.]
	<ul style="list-style-type: none"> Support for global load balancing (GLB). [See Global Load Balancing (GLB).] Support for reactive path rebalancing. [See Reactive Path Rebalancing.] SNMP and telemetry support for PFC, explicit congestion notification (ECN), and CoS ingress packet drops due to ingress port congestion. [See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved, <code>show snmp mib</code>, and Guidelines for gRPC and

Table 1: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
	<p data-bbox="753 352 1370 417">gNMI Sensors (Junos Telemetry Interface). For sensors, see Junos YANG Data Model Explorer.]</p> <ul style="list-style-type: none"> <li data-bbox="719 457 1104 483">• Support for PFC X-ON threshold. <p data-bbox="753 516 1187 541">[See xon (Input Congestion Notification).]</p> <ul style="list-style-type: none"> <li data-bbox="719 581 1390 678">• Extended sFlow monitoring functionality support to export sFlow sample packets through the mgmt_junos interface and non-default virtual routing and forwarding (VRF) WAN ports. <p data-bbox="753 714 1390 774">[See collector, show sflow collector, and System Logging and Routing Instances.]</p> <ul style="list-style-type: none"> <li data-bbox="719 814 1357 875">• Support for configuring per-queue alpha value to limit the buffer each queue can consume from the shared pool. <p data-bbox="753 909 1086 934">[See buffer-dynamic-threshold.]</p> <ul style="list-style-type: none"> <li data-bbox="719 974 1385 1035">• Support for increased global shared buffer pool of up to 147 MB. <p data-bbox="753 1068 1334 1094">[See Configuring Ingress and Egress Dedicated Buffers.]</p> <ul style="list-style-type: none"> <li data-bbox="719 1134 1338 1159">• Support for Inband Flow Analyzer (IFA) 2.0 transit node. <p data-bbox="753 1192 1398 1253">[See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.]</p>

Table 1: QFX5241-32OD Feature Support (*Continued*)

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> • Support for Layer 2 unicast forwarding and VRRP. [See Understanding VRRP.] • Support for IGMP snooping, including: <ul style="list-style-type: none"> • 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) support • Integrated routing and bridging (IRB) support <p>[See IGMP Snooping Overview, Multicast Overview, and Integrated Routing and Bridging.]</p>

Table 1: QFX5241-32OD Feature Support (*Continued*)

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> • 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 (L3) multicast forwarding, including: <ul style="list-style-type: none"> • PIM first-hop router rendezvous point (RP) functionality • Multicast Source Discovery Protocol (MSDP) • Make-before-break (MBB) support for multicast receivers on existing Layer 3 aggregated Ethernet (ae) 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 (PIM SM) • PIM dense mode (PIM DM) • L3 multicast forwarding on integrated routing and bridging (IRB) interfaces functionality, including: <ul style="list-style-type: none"> • IPv4 and IPv6 multicast • IGMPv1, IGMPv2, and IGMPv3 • Multicast Listener Discovery (MLD) versions 1 and 2 • Any-source multicast (ASM) and source-specific multicast (SSM) modes <p>[See Multicast Routing Protocols and PIM Overview.]</p> <ul style="list-style-type: none"> • Support for DHCP stateless relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6. [See DHCP Relay Agent.]

Table 1: QFX5241-32OD Feature Support *(Continued)*

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> Support for sFlow technology. [See Overview of sFlow Technology.] Support for port mirroring and analyzers. The switch can support a maximum of seven port mirroring sessions. [See Understanding Port Mirroring and Analyzers.]
Platform and infrastructure	<ul style="list-style-type: none"> Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See Overview of JET APIs.]
Protection against DDoS attacks	<ul style="list-style-type: none"> Support for configuration and installation of policers at the Packet Forwarding Engine (PFE) level for defense from distributed denial-of-service (DDoS) attacks. By default, the switch provides DDoS protection for many protocols. [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	<ul style="list-style-type: none"> Firewall filter support on Layer 2 and Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions and Configuring Enhanced Egress Firewall Filters.]

Table 1: QFX5241-32OD Feature Support (*Continued*)

Feature	Description
Services applications	<ul style="list-style-type: none"> • Support for generic routing encapsulation (GRE) features, including: <ul style="list-style-type: none"> • GRE tunnels over Gigabit Ethernet, LAG, and VLAN • Tagged subinterfaces • Payload protocol for IPv4 and IPv6 • Delivery protocol for IPv4 • Multicast over GRE tunnels • Tunnel statistics • VRF with GRE • Time-to-live (TTL) <p>[See Generic Routing Encapsulation (GRE).]</p>
Software installation and upgrade	<ul style="list-style-type: none"> • Support for zero-touch provisioning (ZTP) over IPv4 and IPv6 on the management and WAN interfaces. <p>[See Zero Touch Provisioning.]</p>

- **QFX5241-64OD and QFX5241-64QD (QFX Series)**—The Juniper Networks® QFX5241-64OD Switch and Juniper Networks® QFX5241-64QD Switch are fixed-configuration devices with 64 octal small form-factor pluggable (OSFP) or QSFP-DD ports that support speeds up to 800 Gbps. Features such as 51.2 terabits per second (Tbps) throughput and 2-U shallow buffer design make these switches optimal as end-of-row, leaf, or spine devices in IP fabric architectures. The switches support 3000-watt (W) AC and DC power supply units (PSUs) and front-to-back airflow.

To install the QFX5241-64OD and QFX5241-64QD switches and perform initial configuration, routine maintenance, and troubleshooting, see the [QFX5241-64OD and QFX5241-64QD Switches Hardware Guide](#). See [Feature Explorer](#) for the complete list of features for any platform.

Table 2: QFX5241-64OD and QFX5241-64QD Feature Support

Feature	Description
CoS	<ul style="list-style-type: none"> • Support for CoS features on Layer 2 and Layer 3 interfaces, including: <ul style="list-style-type: none"> • 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 for VLAN-tagged traffic. Protocols such as remote direct memory access (RDMA) over Converged Ethernet version 2 (RoCEv2) require DSCP-based PFC at Layer 3 for untagged traffic. • Weighted random early detection (WRED) and explicit congestion notification (ECN) • Telemetry support for CoS queue statistics exported using the sensor <code>/junos/system/linecard/qmon-sw/</code>. <p>[See Traffic Management User Guide (QFX Series Switches and EX4600 Switches).]</p>

Table 2: QFX5241-64OD and QFX5241-64QD Feature Support (*Continued*)

Feature	Description
Features optimized for AI-ML fabrics	<ul style="list-style-type: none"> • Support for priority-based flow control (PFC) watchdog. [See PFC Watchdog and congestion-notification-profile.] • Telemetry support for streaming IPv4 and IPv6 transit statistics using the native resource paths <code>/state/interfaces/interface[name='']/counters/ipv4/</code> and <code>/state/interfaces/interface[name='']/counters/ipv6/</code>. [See Junos YANG Data Model Explorer and route-accounting.] • Support for enabling or disabling dynamic load balancing (DLB). You can use the <code>dynamic-load-balance</code> statement to selectively enable or disable DLB based on <code>rdma-opcode</code> match or any match available in firewall filters. The optimal link is determined based on the modified port load and port queue metrics when DLB is enabled. [See rdma-opcode, dynamic-load-balance-selective, and egress-quantization.] • Support for PFC using DSCPs at Layer 3 for untagged IPv6 traffic. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over Converged Ethernet version 2 (RoCEv2). [See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.] • Support for global load balancing (GLB). [See Global Load Balancing (GLB).] • Support for reactive path rebalancing. [See Reactive Path Rebalancing.] • SNMP and telemetry support for PFC, explicit congestion notification (ECN), and CoS ingress packet drops due to ingress port congestion. [See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved, show snmp mib, and Guidelines for gRPC and

Table 2: QFX5241-64OD and QFX5241-64QD Feature Support (*Continued*)

Feature	Description
	<p>gNMI Sensors (Junos Telemetry Interface). For sensors, see Junos YANG Data Model Explorer.]</p> <ul style="list-style-type: none"> • Support for PFC X-ON threshold. [See xon (Input Congestion Notification).] • Extended sFlow monitoring functionality support to export sFlow sample packets through the mgmt_junos interface and non-default virtual routing and forwarding (VRF) WAN ports. [See collector, show sflow collector, and System Logging and Routing Instances.] • Support for configuring per-queue alpha value to limit the buffer each queue can consume from the shared pool. [See buffer-dynamic-threshold.] • Support for increased global shared buffer pool of up to 147 MB. [See Configuring Ingress and Egress Dedicated Buffers.] • Support for Inband Flow Analyzer (IFA) 2.0 transit node. [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.]

Table 2: QFX5241-64OD and QFX5241-64QD Feature Support (*Continued*)

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> • Support for Layer 2 unicast forwarding and VRRP. [See Understanding VRRP.] • Support for IGMP snooping, including: <ul style="list-style-type: none"> • 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) support • Integrated routing and bridging (IRB) support <p>[See IGMP Snooping Overview, Multicast Overview, and Integrated Routing and Bridging.]</p>

Table 2: QFX5241-64OD and QFX5241-64QD Feature Support *(Continued)*

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> • 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 (L3) multicast forwarding, including: <ul style="list-style-type: none"> • PIM first-hop router rendezvous point (RP) functionality • Multicast Source Discovery Protocol (MSDP) • Make-before-break (MBB) support for multicast receivers on existing Layer 3 aggregated Ethernet (aex) 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 (PIM SM) • PIM dense mode (PIM DM) • L3 multicast forwarding on integrated routing and bridging (IRB) interfaces functionality, including: <ul style="list-style-type: none"> • IPv4 and IPv6 multicast • IGMPv1, IGMPv2, and IGMPv3 • Multicast Listener Discovery (MLD) versions 1 and 2 • Any-source multicast (ASM) and source-specific multicast (SSM) modes [See Multicast Routing Protocols and PIM Overview.] • Support for DHCP stateless relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6. [See DHCP Relay Agent.]

Table 2: QFX5241-64OD and QFX5241-64QD Feature Support (*Continued*)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> Support for sFlow technology. [See Overview of sFlow Technology.] Support for port mirroring and analyzers. These switches can support a maximum of seven port mirroring sessions. [See Understanding Port Mirroring and Analyzers.]
Platform and infrastructure	<ul style="list-style-type: none"> Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See Overview of JET APIs.]
Protection against DDoS attacks	<ul style="list-style-type: none"> Support for configuration and installation of policers at the Packet Forwarding Engine (PFE) level for defense from distributed denial-of-service (DDoS) attacks. By default, these switches provide DDoS protection for many protocols. [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	<ul style="list-style-type: none"> Firewall filter support on Layer 2 and Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions and Configuring Enhanced Egress Firewall Filters.]

Table 2: QFX5241-64OD and QFX5241-64QD Feature Support (*Continued*)

Feature	Description
Services applications	<ul style="list-style-type: none"> • Support for generic routing encapsulation (GRE) features, including: <ul style="list-style-type: none"> • GRE tunnels over Gigabit Ethernet, LAG, and VLAN • Tagged subinterfaces • Payload protocol for IPv4 and IPv6 • Delivery protocol for IPv4 • Multicast over GRE tunnels • Tunnel statistics • VRF with GRE • Time-to-live (TTL) <p>[See Generic Routing Encapsulation (GRE).]</p>
Software installation and upgrade	<ul style="list-style-type: none"> • Support for zero-touch provisioning (ZTP) over IPv4 and IPv6 on the management and WAN interfaces. <p>[See Zero Touch Provisioning.]</p>

Class of Service

- **Support for drop congestion notification (QFX5240-64OD, QFX5240-64QD, QFX5241-64OD, QFX5241-64QD, and QFX5241E-64OD)**—QFX5240 and QFX5241 switches support drop congestion notification (DCN) to enhance congestion management. DCN is a congestion management technique based on packet trimming. Rather than drop a packet when congestion occurs, the device trims the packet's payload, which results in a smaller packet. The device then transmits this smaller packet through a high-priority queue toward its destination. Subsequent hops in the network must recognize DCN-drop marked packets and direct them to high-priority queues as well. End hosts must process the trimmed DCN packets, identify the packets dropped due to congestion, and request retransmission of those lost packets.

[See [Drop Congestion Notification \(DCN\)](#).]

- **DCBX support for PFC (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)**—Data Center Bridging Capability Exchange (DCBX) support for priority-based flow control (PFC).

[See [Understanding DCBX](#) and [priority-flow-control](#).]

- **Dynamic threshold profiles for shared buffer pools (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can tune the buffer allocated per priority group or port based on the dynamic threshold setting, also known as the alpha value. Existing design supports this feature through a global configuration that is applicable to all the ports on the device regardless of the port configuration or properties. Having a global alpha value is not effective when the device has ports operating at various speeds. With this feature, you can create dynamic thresholds per priority group and then associate the dynamic threshold profile on an ingress interface. This configuration creates interface-specific threshold values.

You can configure per priority group dynamic thresholds for each interface to optimize shared buffer allocation and protect lossless traffic during congestion. Egress lossless queues inherit ingress priority-group alpha values.

[See [Dynamic Threshold Profiles for Shared Buffer Pools](#).]

EVPN

- **Hardware-assisted inline BFD for EVPN-VXLAN types 2 and 5 with 3x 100-ms timers (QFX5130-32CD, QFX5130-48C, QFX5130E-32CD, QFX5700, and QFX5700E)**—You can use hardware-assisted inline BFD over VXLAN tunnels for rapid, deterministic failure detection with 100 x 3 millisecond timers on supported platforms. Use IPv4 and IPv6 multihop BFD for Type 2 (L2/L3) with ECMP or multihomed VTEPs. Use Type 5 with ECMP, including pure Type 5 routing instances. To enable hardware-assisted inline BFD, configure bfd on overlay bgp sessions, peer overlays between loopbacks, and apply the specified timers.

[See [Understanding How BFD Detects Network Failures](#).]

Forwarding Options

- **Use the no-queue-pair configuration to exclude queue-pair from the hash calculation (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Previously, the 5 tuple approach was used for hash calculation to identify a flow. But for remote direct memory access (RDMA) traffic, the 5 tuple hashing mechanism wasn't providing the desired entropy. So considering queue-pair into the hash calculation provided better entropy. Currently, queue-pair is by default included in the hash calculation. You can use the no-queue-pair configuration to exclude queue-pair from the hash calculation.

[See [enhanced-hash-key](#).]

- **Unknown unicast drop configuration for VLAN interfaces (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5700E)**—You can enhance network performance and prevent traffic storms by configuring your switch to drop unknown unicast packets. This action prevents the flooding of unicast packets with unknown destination MAC addresses across VLAN interfaces. When you enable this feature, the switch learns and adds the source MAC address to the MAC address table. The switch drops packets with unlearned destination MAC addresses. This approach ensures efficient network resource usage and optimal network performance.

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

Interfaces

- **Monitor pre-FEC BER on optical interfaces (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can assess raw optical signal quality by measuring the pre-forward error correction (FEC) bit error rate (BER), which reflects errors before correction. Use pre-FEC to evaluate the effectiveness of cable performance and error correction. These metrics help you identify impairments early and guide link tuning or maintenance decisions. You collect both metrics in monitoring workflows to support objective path selection, capacity planning, and service assurance. Use the `show interfaces interface-name` command to display the pre-FEC BER. [See [show interfaces](#).]
- **FEC histogram and statistics on optical interfaces (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Use the FEC histogram to evaluate link quality. The histogram shows symbol error corrections per FEC codeword for precise assessment of transmission integrity. Run `show interfaces interface-name` to display FEC counters, corrected and uncorrected codeword counts, pre-FEC BER, and the FEC histogram. This visibility enables you to detect degradation and guide troubleshooting and optimization. [See [show interfaces](#).]
- **Support for native 800 Gbps and SFP28 10/25 G with expanded channelization options for OSFP and QSFP-DD (QFX5241-64OD and QFX5241-64QD)**—Optimize port density and migration by selecting native speeds or channelization on QFX5241-64OD and QFX5241-64QD switches.

On the QFX5241-64OD, you can configure native 800-Gbps speed on the OSFP ports and 10-Gbps speed on the SFP28 ports. You can also channelize OSFP ports into 2x400 Gbps or 8x100 Gbps.

On the QFX5241-64QD:

- Configure the following native speeds on the QSFP-DD ports:
 - 800 Gbps
 - 400 Gbps
 - 100 Gbps
- Use 10-Gbps speed on the SFP28 ports.

You can also channelize the QSFP-DD ports into the following speeds:

- 2x400 Gbps
- 4x100 Gbps
- 8x100 Gbps
- 8x50 Gbps

Do not mix 8x100 Gbps or 8x50 Gbps on odd and even ports. See [Port Speed on QFX Series Switches](#). Also, Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module.

- **Port speed and channelization (QFX5241-32OD)**— You can tailor port speeds and channelize interfaces on QFX5241-32OD switches. On the QFX5241-32OD, you can configure native 800-Gbps speed on the OSFP ports and 10-Gbps speed on the SFP28 ports.

You can also channelize the OSFP ports into the following speeds:

- 1x800 Gbps
- 2x400 Gbps
- 4x200 Gbps

Even-numbered ports support 8x100G or 8x50G; the next odd-numbered port supports up to 2x. [See [Port Speed on QFX Switches](#).]

Junos Telemetry

- **Export packet capture statistics to external collectors using Junos Telemetry (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, and QFX5700)**—The packet capture feature records the first configured number of host-bound packets on each physical interface and exports them to an external collector over the Junos Telemetry infrastructure. You can use this data to debug and fix network or performance issues. Subscribe to the packet-capture sensor at `/junos/system/linecard/packet-capture`. The device captures 50 ingress packets when an interface transitions from the DOWN state to the UP state. The data is encoded in Google protocol buffer (GPB) format and streamed over gRPC with SSL encryption. To enable packet capture, configure `edit system packet-forwarding-options packet-capture packet-capture-enable`.

Use the `show agent sensors` command to view the packet capture sensor information.

View the complete list of supported sensors in the [Junos YANG Data Model Explorer](#).

Layer 2 VPN

- **Support for rewriting the L2PT destination MAC address (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5700E)**—In Layer 2 Protocol Tunneling (L2PT), the device rewrites the original multicast destination MAC address of the packet. The packet then travels across the provider network transparently to the other end of the tunnel, where the destination device restores the original multicast destination MAC address. By default, the device rewrites the multicast destination MAC address with the predefined multicast tunneling MAC address 01:00:0C:CD:CD:D0 in the MAC table. You can optionally specify a different multicast MAC address.

To specify the MAC address, use the `tunnel-destination-mac mac-address` statement at the `[edit protocols layer2-control layer2-control mac-rewrite]` hierarchy level. You can set any non-reserved multicast MAC address.

[See [mac-rewrite](#) and [Layer 2 Protocol Tunneling \(L2PT\)](#).]

Network Management and Monitoring

- **Timestamp option for tap-aggregation packets (QFX5220-32CD and QFX5220-128C)**—High-performance data analysis depends on the accuracy of its underlying data; one accuracy enhancement is a timestamp inserted into a tap-aggregation packet that shows exactly when the packet was captured.

Configure the tap-aggregation feature to insert a timestamp in packets at data capture, before the packets are sent to the tool ports for analysis. Before the timestamp is inserted, configure the PTP reference clock on the tap-aggregation switch and confirm that PTP is running.. Your tap-aggregation switch must also synchronize the PTP FPGA's recovered time-of-day with the system chip's time-of-day. Enable timestamping per interface with

```
[edit] user@switch# set interfaces interface-name timestamp ingress
```

- **Ingress ACL UDF filtering function on tap ports on TAP-aggregation switches (QFX5130-32CD, QFX5130-48C, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, and QFX5700)**—Use an ingress access control list (ACL) with user-defined field (UDFs) on tap interfaces to direct traffic to tool interfaces on a tap-aggregation switch.

If an ACL match conflicts with a tap-aggregation rule, the ACL match takes precedence.

Configure the tap-aggregation interfaces under the `[edit interfaces]` hierarchy level as follows:

- Add an interface to a tap group with

```
[edit]
```

```
user@switch# set interfaces interface-name unit 0 mode tap group tap-group-name
```


- Add an interface to a tool group with

[edit]

```
user@switch# set interfaces interface-name unit 0 mode tool group tool-group-name
```

- **Dropped-packet notification (QFX5240-64OD and QFX5240-64QD)**—Packet drops commonly occur on network devices, and debugging can be complex and can result in high mean times to recovery (MTTRs). Although the packet-processing pipeline provides some drop counters, they are insufficient for debugging complex packet-drop issues. Use dropped-packet notification (also called mirror on drop, or MoD) to debug drops in real time. This feature monitors packet drops caused due to processing in the ingress and egress pipelines and due to congestion in the memory management unit (MMU). This feature is stateless and flow unaware. Configure most of the feature at the [edit forwarding-options mirror-profile] hierarchy level.

Routing Protocols

- **Enhancing Traffic Distribution with Weighted Packet Spray (QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, and QFX5241-64QD)**—Weighted Packet Spray (WPS) is a load balancing technique that enhances traffic distribution across network fabrics by assigning weights to packets and rebalancing flows based on link bandwidth. Operating at the route level, WPS improves efficiency and resilience during link failures and works alongside Weighted ECMP to optimize traffic paths. Firewall filters with the weighted-packet-spray option allow selective application to specific flows, while other flows continue using flow-based Weighted ECMP. For BGP-based load balancing, all paths of a given route must advertise link bandwidth; if any path lacks this attribute, that route alone falls back to equal-cost multipath (ECMP). Consistent deployment of WPS across network tiers reduces packet loss, maximizes bandwidth utilization, and ensures stable load balancing.

[See [Weighted Packet Spray](#).]

- **GLB multi-link support on IP Fabric (QFX5240-64OD and QFX5240-64QD)**—We are extending Global Load Balancing (GLB) to support multiple paths between spine and top-of-rack switches on a 3-stage CLOS IP fabric.

To enable this feature, include the `glb-multilink-mode max-val/avg-val` statement at the [edit forwarding-options enhanced-hash-key] hierarchy level. By default, the spine advertises the average quality of all links. Make sure you enable GLB globally at the [edit protocols bgp] hierarchy level.

[See [Configure GLB on 3-CLOS IP Fabric with Multilinks](#).]

- **Support for 256-way ECMP (QFX5130-32CD, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Use this feature to increase the number of direct Border Gateway Protocol (BGP) peer connections, improve latency, and optimize data flow by configuring up to 256 equal-cost multipath (ECMP) next hops for external BGP peers.

[See [Example: Load Balancing BGP Traffic.](#)]

Additional Features

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

- **EVPN-VXLAN Layer 3 gateway** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD). Support includes:
 - Layer 3 (L3) VXLAN gateway in edge-routed bridging (ERB) fabric
 - Up to 256 integrated routing and bridging (IRB) enabled VLANs
 - Layer 3 underlay that supports IRB and link aggregation groups (LAGs)
 - ECMP in the underlay
 - IPv4 and IPv6 virtual gateway MAC address support for IRB interfaces
 - In-service software upgrade (ISSU) for L3 gateway functionality

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#), [Example: Configuring an EVPN-VXLAN Edge-Routed Bridging Fabric with a Virtual Gateway](#), [Understanding the MAC Addresses For a Default Virtual Gateway in an EVPN-VXLAN or EVPN-MPLS Overlay Network](#), and [IP Fabric Underlay Network Design and Implementation.](#)]

- **EVPN-VXLAN Layer 2 gateway** (QFX5240-64OD and QFX5240-64QD). Support includes:
 - Up to 508 VXLAN-enabled VLANs
 - Layer 3 underlay with integrated routing and bridging (IRB) and link aggregation groups (LAGs)
 - ECMP in the underlay
 - VTEP ingress and egress counters
 - BUM traffic forwarding by ingress replication only
 - MAC move functionality
 - Core isolation support as a control-plane feature
 - `no-mac-ip-learning` command to disable `mac-ip` learning
 - MAC move limit enforcement with port disable and VLAN member disable actions
 - In-service software upgrade (ISSU) for Layer 3 gateway functionality

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#), [IP Fabric Underlay Network Design and Implementation](#), [Overview of Hierarchical ECMP Groups on QFX5200 Switches](#), [NSR and Unified ISSU Support for EVPN](#), and [irb-symmetric-routing](#).]

- **EVPN-VXLAN Type 5 route** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD). Support includes:
 - ECMP in underlay and overlay
 - Layer 3 underlay with integrated routing and bridging (IRB) and link aggregation groups (LAGs)
 - IPv6 user traffic support
 - In-service software upgrade (ISSU) support
 - Type 5 seamless stitching
 - Symmetric IRB
 - Type 2 and Type 5 route coexistence

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#), [IP Fabric Underlay Network Design and Implementation](#), [overlay-ecmp](#), [Understanding EVPN Pure Type 5 Routes](#), [EVPN Type 2 and Type 5 Route Coexistence with EVPN-VXLAN](#), [NSR and Unified ISSU Support for EVPN](#), and [irb-symmetric-routing](#).]

- **Support for CoS and firewall filtering and policing on EVPN-VXLAN network** (QFX5240-64OD and QFX5240-64QD).

[See [CoS Support on EVPN VXLANs](#) and [Firewall Filter Match Conditions and Actions \(QFX and EX Series Switches\)](#).]

- **Support for sFlow technology on EVPN-VXLAN** (QFX5240-64OD and QFX5240-64QD).

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

- **Support for port mirroring and analyzers on EVPN-VXLAN** (QFX5240-64OD and QFX5240-64QD).

[See [Port Mirroring and Analyzers in an EVPN-VXLAN Environment](#).]

- **Fast reroute for egress link protection (ELP) in EVPN-VXLAN multihoming environments** (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5700E).

[See [Fast Reroute for Egress Link Protection with EVPN-VXLAN Multihoming](#) and [reroute-address](#).]

- **Simplified configuration for ESI LAGs with EVPN dual homing (EZ-LAG)** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD).

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

- **Separate firmware installation packages** (QFX5130, QFX5220, QFX5230, QFX5240, QFX5700). You can manage firmware upgrades using standalone firmware installation packages. The names for these packages begin with the prefix `jfirmware-junos-evo-install*`. You add a firmware package to the system with the `request system software add` command. After you've added the firmware package to your system, you update the firmware for a hardware component using the `request system firmware upgrade` command.

[See [Upgrade Firmware on Junos OS Evolved Devices](#), [Junos OS Evolved Installation Packages](#), [request system software add \(Junos OS Evolved\)](#), [request system firmware upgrade \(Junos OS Evolved\)](#), and [show system firmware \(Junos OS Evolved\)](#).]

- **Supported transceivers, optical interfaces, and DAC cables** (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5220-32CD, QFX5240-64OD, QFX5240-64QD, QFX5230-64CD, and QFX5700). Select your product in the [Hardware Compatibility Tool](#) to view supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update the HCT and provide the first supported release information when the optical transceiver becomes available.
- **Symmetric IRB for IPv6 underlay with EVPN Type 2 MAC-IP routes for intersubnet routing** (QFX5130-32CD and QFX5130-48C). You can now use symmetrical integrated routing and bridging (IRB) to forward traffic consistently across all provider edge (PEs) devices. The device carries EVPN Type 2 MAC-IP routes across an IPv6 underlay that provides an ECMP transport for scalable, multitenant Layer 2 (L2) and Layer 3 (L3) connectivity. Each routing instance uses the `irb-symmetric-routing vni` statement under the `protocols evpn` hierarchy to maintain symmetrical bridging and routing.

[See [EVPN-VXLAN with an IPv6 Underlay](#) and [Symmetric Integrated Routing and Bridging with EVPN Type 2 Routes in EVPN-VXLAN Fabrics](#).]

- **Support for zero-touch provisioning (ZTP)** (QFX5250-64OE-L and QFX5250-64OE-AO).
[See [Zero Touch Provisioning](#).]
- **Support for secure zero-touch provisioning (SZTP)** (QFX5250-64OE-L and QFX5250-64OE-AO).
[See [Secure Zero Touch Provisioning](#).]
- **Support for wildcard mask match condition for source-address/destination-address match conditions for inet6 address family** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD).

[See [Understanding Firewall Filter Match Conditions](#) and [IPv6 Wildcard Mask Match Conditions](#).]

What's Changed

There are no changes in behavior and syntax in this release for QFX Series switches.

Known Limitations

There are no known limitations in hardware or software in Junos OS Evolved Release 25.2X100-D10 for QFX Series switches.

Open Issues

IN THIS SECTION

- [MACsec \(Media Access Control Security\) | 27](#)

Learn about open issues in Junos OS Evolved Release 25.2X100-D10 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.

MACsec (Media Access Control Security)

- QFX5130-48C/QFX5130-48CM MACsec: On MACsec enabled interfaces, after evo-pfemamd restart with cli restart evo-pfemamd, ping might not work sometimes. [PR1891444](#)

Licensing

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

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.

- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information about the list of supported products, see [Juniper Flex Program](#).

Finding More Information

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

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

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

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

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

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



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

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

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

Requesting Technical Support

IN THIS SECTION

- Self-Help Online Tools and Resources | 29
- Creating a Service Request with JTAC | 30

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <https://www.juniper.net/content/dam/www/assets/resource-guides/us/en/jtac-user-guide.pdf>.
- Product warranties—For product warranty information, visit <https://support.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://support.juniper.net/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://supportportal.juniper.net/s/knowledge>

- Download the latest versions of software and review release notes: <https://support.juniper.net/support/downloads/>
- Search technical bulletins for relevant hardware and software notifications: <https://supportportal.juniper.net/s/knowledge>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://supportportal.juniper.net/>

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

Creating a Service Request with JTAC

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

- Visit <https://support.juniper.net/support/requesting-support/>
- 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 December 2025—Revision 1, Junos OS Evolved Release 25.2X100-D10

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice. Copyright © 2025 Juniper Networks, Inc. All rights reserved.