

# Release Notes

Published  
2025-12-22

## Junos OS Evolved Release 23.4X100-D40

### Introduction

Use these release notes to find new features, software limitations, and open issues for Junos OS Evolved Release 23.4X100-D40.

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



**NOTE:** Junos OS Evolved 23.4X100-D40 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-OD or 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 | 2

Hardware   2
Chassis   22
Class of Service   23
EVPN   24
Forwarding Options   25
High Availability   25
Interfaces   25
Junos Telemetry Interface   25
Network Management and Monitoring   26
Routing Protocols   27
Software Installation and Upgrade   28
Additional Features   28

### What's Changed | 29

### Known Limitations | 29

### Open Issues | 30

### Licensing | 31

### Finding More Information | 32

### Requesting Technical Support | 32

### Revision History | 34

# Junos OS Evolved Release Notes for QFX Series

## IN THIS SECTION

- [What's New | 2](#)
- [What's Changed | 29](#)
- [Known Limitations | 29](#)
- [Open Issues | 30](#)

These release notes accompany Junos OS Evolved Release 23.4X100-D40 for QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220-32CD, QFX5220-128C, QFX5230-64CD, QFX5240-OD, QFX5240-QD, QFX5700, and QFX5700E switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.



**NOTE:** We recommend Junos OS Evolved 23.4X100-D43 on the following QFX platforms:

- QFX5240-64OD
- QFX5240-64QD
- QFX5241-32OD
- QFX5241-64OD
- QFX5241-64QD

For more information, see the [Junos OS Evolved 23.4X100-D43 Software Release Notification](#).

## What's New

### IN THIS SECTION

- [Hardware | 2](#)
- [Chassis | 22](#)
- [Class of Service | 23](#)
- [EVPN | 24](#)
- [Forwarding Options | 25](#)
- [High Availability | 25](#)
- [Interfaces | 25](#)
- [Junos Telemetry Interface | 25](#)
- [Network Management and Monitoring | 26](#)
- [Routing Protocols | 27](#)
- [Software Installation and Upgrade | 28](#)
- [Additional Features | 28](#)

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 gigabits per second (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"> <li>• Support for CoS features on Layer 2 and Layer 3 interfaces, including: <ul style="list-style-type: none"> <li>• IPv4 and IPv6 unicast routing</li> <li>• Classification and rewrite rules for Differentiated Services code point (DSCP) and IEEE-802.1p</li> <li>• Port scheduling</li> <li>• Shared buffer</li> </ul> </li> <li>• 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.</li> <li>• Weighted random early detection (WRED) and explicit congestion notification (ECN)</li> <li>• Telemetry support for CoS queue statistics exported using the sensor <code>/junos/system/linecard/qmon-sw/</code>.</li> </ul> <p>[See <a href="#">Traffic Management User Guide (QFX Series Switches and EX4600 Switches)</a>.]</p>

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

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

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

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

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

Feature	Description
	<p>[See <a href="#">Overview of sFlow Technology</a>.]</p> <ul style="list-style-type: none"><li>• Support for port mirroring and analyzers on EVPN-VXLAN.</li></ul> <p>[See <a href="#">Port Mirroring and Analyzers in an EVPN-VXLAN Environment</a>.]</p> <ul style="list-style-type: none"><li>• 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.</li></ul> <p>[See <a href="#">Understanding Unified ISSU</a>.]</p>



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

Feature	Description
Features optimized for AI-ML fabrics	<ul style="list-style-type: none"> <li>Support for priority-based flow control (PFC) watchdog. [See <a href="#">PFC Watchdog</a> and <a href="#">congestion-notification-profile</a>.]</li> </ul>
	<ul style="list-style-type: none"> <li>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 <a href="#">Junos YANG Data Model Explorer</a> and <a href="#">route-accounting</a>.]</li> </ul>
	<ul style="list-style-type: none"> <li>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 <a href="#">rdma-opcode</a>, <a href="#">dynamic-load-balance-selective</a>, and <a href="#">egress-quantization</a>.]</li> </ul>
	<ul style="list-style-type: none"> <li>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 <a href="#">Understanding PFC Using DSCP at Layer 3 for Untagged Traffic</a>.]</li> </ul>
	<ul style="list-style-type: none"> <li>Support for global load balancing (GLB). [See <a href="#">Global Load Balancing (GLB)</a>.]</li> <li>Support for reactive path rebalancing. [See <a href="#">Reactive Path Rebalancing</a>.]</li> <li>SNMP and telemetry support for PFC, explicit congestion notification (ECN), and CoS ingress packet drops due to ingress port congestion. [See <a href="#">SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved</a>, <a href="#">show snmp mib</a>, and <a href="#">Guidelines for gRPC and</a></li> </ul>

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

Feature	Description
	<p data-bbox="753 352 1370 422"><a href="#">gNMI Sensors (Junos Telemetry Interface)</a>. For sensors, see <a href="#">Junos YANG Data Model Explorer</a>.]</p> <ul style="list-style-type: none"> <li data-bbox="719 457 1094 485">• Support for PFC XON threshold.</li> </ul> <p data-bbox="753 516 1187 543">[See <a href="#">xon (Input Congestion Notification)</a>.]</p> <ul style="list-style-type: none"> <li data-bbox="719 579 1386 680">• Extended sFlow monitoring functionality support to export sFlow sample packets through the mgmt_junos interface and nondefault virtual routing and forwarding (VRF) WAN ports.</li> </ul> <p data-bbox="753 711 1386 779">[See <a href="#">collector</a>, <a href="#">show sflow collector</a>, and <a href="#">System Logging and Routing Instances</a>.]</p> <ul style="list-style-type: none"> <li data-bbox="719 814 1354 877">• Support for configuring per-queue alpha value to limit the buffer each queue can consume from the shared pool.</li> </ul> <p data-bbox="753 909 1086 936">[See <a href="#">buffer-dynamic-threshold</a>.]</p> <ul style="list-style-type: none"> <li data-bbox="719 972 1386 1035">• Support for increased global shared buffer pool of up to 147 MB.</li> </ul> <p data-bbox="753 1066 1333 1094">[See <a href="#">Configuring Ingress and Egress Dedicated Buffers</a>.]</p> <ul style="list-style-type: none"> <li data-bbox="719 1129 1338 1157">• Support for Inband Flow Analyzer (IFA) 2.0 transit node.</li> </ul> <p data-bbox="753 1188 1398 1255">[See <a href="#">Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring</a>.]</p>

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

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> <li>• Support for Layer 2 unicast forwarding and VRRP. [See <a href="#">Understanding VRRP</a>.]</li> <li>• Support for IGMP snooping, including: <ul style="list-style-type: none"> <li>• IGMPv1, IGMPv2, and IGMPv3</li> <li>• IGMP proxy</li> <li>• IGMP querier at Layer 2</li> <li>• Any-source multicast (ASM) and source-specific multicast (SSM) modes</li> <li>• Virtual router (VRF-lite) support</li> <li>• Integrated routing and bridging (IRB) support</li> </ul> </li> </ul> <p>[See <a href="#">IGMP Snooping Overview</a>, <a href="#">Multicast Overview</a>, and <a href="#">Integrated Routing and Bridging</a>.]</p>

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

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> <li>• Support for Layer 3 unicast forwarding and generic routing encapsulation (GRE) tunneling. We support both IPv4 and IPv6 unicast routing.  [See <a href="#">Generic Routing Encapsulation (GRE)</a>.]</li> <li>• Support for Layer 3 (L3) multicast forwarding, including: <ul style="list-style-type: none"> <li>• PIM first-hop router rendezvous point (RP) functionality</li> <li>• Multicast Source Discovery Protocol (MSDP)</li> <li>• 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.</li> <li>• PIM source-specific multicast (SSM)</li> <li>• PIM sparse mode (PIM SM)</li> <li>• PIM dense mode (PIM DM)</li> </ul> </li> <li>• L3 multicast forwarding on integrated routing and bridging (IRB) interfaces functionality, including: <ul style="list-style-type: none"> <li>• IPv4 and IPv6 multicast</li> <li>• IGMPv1, IGMPv2, and IGMPv3</li> <li>• Multicast Listener Discovery (MLD) versions 1 and 2</li> <li>• Any-source multicast (ASM) and source-specific multicast (SSM) modes</li> </ul> </li> </ul> <p>[See <a href="#">Multicast Routing Protocols</a> and <a href="#">PIM Overview</a>.]</p> <ul style="list-style-type: none"> <li>• Support for DHCP stateless relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6.  [See <a href="#">DHCP Relay Agent</a>.]</li> </ul>

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

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> <li>• Support for sFlow technology. [See <a href="#">Overview of sFlow Technology</a>.]</li> <li>• Support for port mirroring and analyzers. The switch supports a maximum of seven port mirroring sessions. [See <a href="#">Understanding Port Mirroring and Analyzers</a>.]</li> </ul>
Platform and infrastructure	<ul style="list-style-type: none"> <li>• Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See <a href="#">Overview of JET APIs</a>.]</li> </ul>
Protection against DDoS attacks	<ul style="list-style-type: none"> <li>• 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, DDoS protection is enabled for many protocols on the switch. [See <a href="#">Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers</a>, <a href="#">show ddos-protection statistics</a>, and <a href="#">show ddos-protection version</a>.]</li> </ul>
Routing policy and firewall filters	<ul style="list-style-type: none"> <li>• Firewall filter support on Layer 2 and Layer 3 interfaces. [See <a href="#">Firewall Filter Match Conditions and Actions</a> and <a href="#">Configuring Enhanced Egress Firewall Filters</a>.]</li> </ul>

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

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

- **QFX5241-64OD and QFX5241-64QD (QFX Series)**—The QFX5241-64OD and QFX5241-64QD switches are fixed-configuration devices with 64 octal small form-factor pluggable (OSFP) or QSFP-DD ports that support speeds of up to 800 gigabits per second (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"> <li>• Support for CoS features on Layer 2 and Layer 3 interfaces, including:             <ul style="list-style-type: none"> <li>• IPv4 and IPv6 unicast routing</li> <li>• Classification and rewrite rules for Differentiated Services code point (DSCP) and IEEE-802.1p</li> <li>• Port scheduling</li> <li>• Shared buffer</li> </ul> </li> <li>• 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 Differentiated Services code point (DSCP)-based PFC at Layer 3 for untagged traffic.</li> <li>• Weighted random early detection (WRED) and explicit congestion notification (ECN)</li> <li>• Telemetry support for CoS queue statistics exported using the sensor <code>/junos/system/linecard/qmon-sw/</code>.</li> </ul> <p>[See <a href="#">Traffic Management User Guide (QFX Series Switches and EX4600 Switches)</a>.]</p>

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

Feature	Description
EVPN	<ul style="list-style-type: none"> <li>• Support for Layer 2 gateway and Address Resolution Protocol (ARP) suppression on EVPN-VXLAN. [See <a href="#">Understanding EVPN with VXLAN Data Plane Encapsulation, EVPN Proxy ARP and ARP Suppression, and overlay (Packet Forwarding Options)</a>.]</li> <li>• Support for CoS and firewall filtering and policing on EVPN-VXLAN. [See <a href="#">CoS Support on EVPN VXLANs</a> and <a href="#">Firewall Filter Match Conditions and Actions (QFX and EX Series Switches)</a>.]</li> <li>• Support for Wake-on-LAN (WOL) targeted broadcast on EVPN-VXLAN. [See <a href="#">Targeted Broadcast</a> and <a href="#">targeted-broadcast</a>.]</li> <li>• Support for EVPN-VXLAN Layer 2 gateway, including: <ul style="list-style-type: none"> <li>• Multihoming</li> <li>• ARP suppression</li> <li>• Layer 3 IPv4 underlay with integrated routing and bridging (IRB) and LAG</li> <li>• Core isolation</li> <li>• Broadcast, unknown unicast, and multicast (BUM) traffic forwarding by ingress replication only</li> <li>• MAC move limits</li> </ul> [See <a href="#">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</a>.]</li> </ul>



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

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

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

Feature	Description
	<ul style="list-style-type: none"><li>• Support for sFlow technology on EVPN-VXLAN. [See <a href="#">Overview of sFlow Technology</a>.]</li><li>• Support for port mirroring and analyzers on EVPN-VXLAN. [See <a href="#">Port Mirroring and Analyzers in an EVPN-VXLAN Environment</a>.]</li><li>• 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. [See <a href="#">Understanding Unified ISSU</a>.]</li></ul>

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

Feature	Description
Features optimized for AI-ML fabrics	<ul style="list-style-type: none"> <li>• Support for priority-based flow control (PFC) watchdog. [See <a href="#">PFC Watchdog</a> and <a href="#">congestion-notification-profile</a>.]</li> <li>• 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 <a href="#">Junos YANG Data Model Explorer</a> and <a href="#">route-accounting</a>.]</li> <li>• 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 <a href="#">rdma-opcode</a>, <a href="#">dynamic-load-balance-selective</a>, and <a href="#">egress-quantization</a>.]</li> <li>• 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 <a href="#">Understanding PFC Using DSCP at Layer 3 for Untagged Traffic</a>.]</li> <li>• Support for global load balancing (GLB). [See <a href="#">Global Load Balancing (GLB)</a>.]</li> <li>• Support for reactive path rebalancing. [See <a href="#">Reactive Path Rebalancing</a>.]</li> <li>• SNMP and telemetry support for PFC, explicit congestion notification (ECN), and CoS ingress packet drops due to ingress port congestion. [See <a href="#">SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved</a>, <code>show snmp mib</code>, and <a href="#">Guidelines for gRPC and</a></li> </ul>

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

Feature	Description
	<p><a href="#">gNMI Sensors (Junos Telemetry Interface)</a>. For sensors, see <a href="#">Junos YANG Data Model Explorer</a>.]</p> <ul style="list-style-type: none"> <li>• Support for PFC XON threshold. [See <a href="#">xon (Input Congestion Notification)</a>.]</li> <li>• 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 <a href="#">collector</a>, <a href="#">show sflow collector</a>, and <a href="#">System Logging and Routing Instances</a>.]</li> <li>• Support for configuring per-queue alpha value to limit the buffer each queue can consume from the shared pool. [See <a href="#">buffer-dynamic-threshold</a>.]</li> <li>• Support for increased global shared buffer pool of up to 147 MB. [See <a href="#">Configuring Ingress and Egress Dedicated Buffers</a>.]</li> <li>• Support for Inband Flow Analyzer (IFA) 2.0 transit node. [See <a href="#">Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring</a>.]</li> </ul>

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

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> <li>• Support for Layer 2 unicast forwarding and VRRP. [See <a href="#">Understanding VRRP</a>.]</li> <li>• Support for IGMP snooping, including:             <ul style="list-style-type: none"> <li>• IGMPv1, IGMPv2, and IGMPv3</li> <li>• IGMP proxy</li> <li>• IGMP querier at Layer 2</li> <li>• Any-source multicast (ASM) and source-specific multicast (SSM) modes</li> <li>• Virtual router (VRF-lite) support</li> <li>• Integrated routing and bridging (IRB) support</li> </ul> </li> </ul> <p>[See <a href="#">IGMP Snooping Overview</a>, <a href="#">Multicast Overview</a>, and <a href="#">Integrated Routing and Bridging</a>.]</p>

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

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> <li>• Support for Layer 3 unicast forwarding and generic routing encapsulation (GRE) tunneling. We support both IPv4 and IPv6 unicast routing. [See <a href="#">Generic Routing Encapsulation (GRE)</a>.]</li> <li>• Support for Layer 3 (L3) multicast forwarding, including: <ul style="list-style-type: none"> <li>• PIM first-hop router rendezvous point (RP) functionality</li> <li>• Multicast Source Discovery Protocol (MSDP)</li> <li>• 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.</li> <li>• PIM source-specific multicast (SSM)</li> <li>• PIM sparse mode (PIM SM)</li> <li>• PIM dense mode (PIM DM)</li> </ul> </li> <li>• L3 multicast forwarding on integrated routing and bridging (IRB) interfaces functionality, including: <ul style="list-style-type: none"> <li>• IPv4 and IPv6 multicast</li> <li>• IGMPv1, IGMPv2, and IGMPv3</li> <li>• Multicast Listener Discovery (MLD) versions 1 and 2</li> <li>• Any-source multicast (ASM) and source-specific multicast (SSM) modes</li> </ul> </li> <li>[See <a href="#">Multicast Routing Protocols</a> and <a href="#">PIM Overview</a>.]</li> <li>• Support for DHCP stateless relay on IRB interfaces and bridge domains. Support includes DHCPv4 and DHCPv6. [See <a href="#">DHCP Relay Agent</a>.]</li> </ul>

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

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> <li>Support for sFlow technology. [See <a href="#">Overview of sFlow Technology</a>.]</li> <li>Support for port mirroring and analyzers. The QFX5240-64OD and QFX5240-64QD switches can support a maximum of seven port mirroring sessions. [See <a href="#">Understanding Port Mirroring and Analyzers</a>.]</li> </ul>
Platform and infrastructure	<ul style="list-style-type: none"> <li>Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See <a href="#">Overview of JET APIs</a>.]</li> </ul>
Protection against DDoS attacks	<ul style="list-style-type: none"> <li>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, DDoS protection is enabled for many protocols on these switches. [See <a href="#">Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers</a>, <a href="#">show ddos-protection statistics</a>, and <a href="#">show ddos-protection version</a>.]</li> </ul>
Routing policy and firewall filters	<ul style="list-style-type: none"> <li>Firewall filter support on Layer 2 and Layer 3 interfaces. [See <a href="#">Firewall Filter Match Conditions and Actions</a> and <a href="#">Configuring Enhanced Egress Firewall Filters</a>.]</li> </ul>

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

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

## Chassis

- **Resiliency support (QFX5241-32OD)**—We support resiliency for platform components, which enables the system to monitor component health, alert you of errors, and take appropriate action to restore normal operation based on error severity.

[See [Resiliency](#).]

- **Chassis management support (QFX5241-32OD)**—We provide support for built-in Routing Engines, Control Boards, power supply units (PSUs), fan trays, PICs, and flexible PIC concentrators (FPCs) on switches with AC PSUs. The QFX5241-32OD Switch supports all features relevant to fixed-form-factor switches. It also support these redundancy features:
  - N+1 fan redundancy. The fans are field replaceable units (FRUs). You can manage the fan speed according to the environment monitoring policy.



- Two PSUs with 1+1 redundancy.

You can manage the switch and its components such as sensors and interfaces, by using existing CLI commands.

[See [temperature sensor](#) and [Configure Port Speed \(\)](#).]

## Class of Service

- **Drop congestion notification (QFX5240-64OD, QFX5240-64QD, QFX5241-64OD, and QFX5241-64QD)**—QFX5240 and QFX5241 switches can function as transit switches for drop congestion notification (DCN). DCN is a congestion management technique based on packet trimming. Rather than dropping a packet when congestion occurs, the packet's payload is trimmed, resulting in a smaller packet that is then transmitted through a high-priority queue toward its destination. Subsequent hops in the network recognize DCN-marked packets and direct them to high-priority queues. End hosts must process the trimmed DCN packets, identify the packets dropped due to congestion, and request retransmission of those lost packets. The original sender host, upon receiving the DCN feedback, immediately re-transmits the dropped packet and adjusts the flow rate.

DCN works only for UDP unicast traffic. To enable DCN on the switch, set the UDP port number (`udp-port port-number`) and a strict-high priority forwarding class (`forwarding-class class-name`) at the `[edit class-of-service drop-congestion-notification]` hierarchy level. The L4 UDP port number is set to the DCN protocol number so that the switch can identify the DCN packets.

Then set `drop-congestion-notification` on an ingress interface at the `[edit class-of-service interface ingress-interface-name]` hierarchy level so that the interface can detect the DCN-marked packets.

- **Per-interface ingress shared buffer pool allocation (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—You can tune the buffer allocated to each priority group or interface based on the dynamic threshold value (also referred to as the alpha value).

You can set the ingress dynamic-threshold value on a global basis at the `[edit class-of-service shared-buffer ingress buffer-partition (lossless, lossless-headroom, lossy)]` hierarchy level. A global alpha value is useful in certain cases, but it is not effective when the switch has ports operating at different speeds.

For switches with ports of different speeds, you can define a dynamic threshold profile with an alpha value for each priority group and then associate the profile with an ingress interface. This configuration enables interface-specific dynamic threshold values.

To enable per-interface dynamic threshold values, define a dynamic threshold profile at the `[edit class-of-service]` hierarchy level. For example:

```
set class-of-service dynamic-threshold-profile profile-name ingress priority-group PG-number dynamic-threshold threshold-value
```

Then assign the dynamic threshold profile to an ingress interface. For example:

```
set class-of-service interfaces interface-name dynamic-threshold-profile profile-name
```

An interface-specific dynamic threshold value overrides a defined global dynamic threshold value or the default dynamic threshold value.

## EVPN

- **Inline BFD with EVPN-VXLAN for Type 2 and Type 5 tunnel repair (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5700, and QFX5700E)**—We support inline BFD with EVPN-VXLAN for faster recovery of PE to PE network failures.

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

- **Override global AS number for auto-derived route targets in EVPN instances (QFX5130-32CD)**—By default, EVPN provider edge devices use the global autonomous system (AS) number to auto-derive route targets and configure the related route target import and export policies in EVPN instances (EVIs). If required, you can override this default behavior to instead specify a common AS number to use for auto-derived route targets across EVPN peer devices. With this feature, you can ensure consistent route advertisements and import policies in EVPN fabrics that use external BGP (eBGP) underlay and overlay peering.

Use the `import-as` and `export-as` options to assign an AS number on all EVPN peer devices for auto-derived route target import and export policies. Alternatively, you can use the `as-num` option to assign the same AS number for both import and export policies.

- For default-switch instances:

```
set switch-options vrf-target auto import-as asnum vni-list <all | [vni-1 vni-2 ...]>
set switch-options vrf-target auto export-as asnum vni-list <all | [vni-1 vni-2 ...]>
set switch-options vrf-target auto as-num asnum vni-list <all | [vni-1 vni-2 ...]>
```

- For mac-vrf instances:

```
set routing-instances name vrf-target auto import-as asnum vni-list <all | [vni-1 vni-2 ...]>
set routing-instances name vrf-target auto export-as asnum vni-list <all | [vni-1 vni-2 ...]>
set routing-instances name vrf-target auto as-num asnum vni-list <all | [vni-1 vni-2 ...]>
```

[See [Auto-derived Route Targets](#) and [vrf-target](#).]

## Forwarding Options

**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 unnecessary 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 resources usage and optimal network performance.

## High Availability

- **Hardware-assisted inline BFD support (QFX5130-48C and QFX5230-64CD)**—You can leverage hardware-assisted inline BFD to monitor and maintain network connectivity. With hardware-assisted inline BFD, the firmware handles most of the BFD protocol processing, leading to faster processing times compared to traditional inline BFD. Use the `routing-options ppm inline-processing-enable` configuration statement to enable hardware-assisted inline BFD.

[See [Hardware-Assisted Inline BFD](#).]

## Interfaces

- **Port configurations (QFX5241-32OD)**—We offer native 800 G on octal small form-factor pluggable (OSFP) ports and native 10 G on SFP28 ports. Additionally, you can channelize OSFP ports into various configurations such as 2x400G, 8x100G, 8x50G, 4x200G, 1x400G, 2x200G, and 4x100G. Use these port configurations to optimize network performance and scalability according to your requirements.
- **Hold-timer value configuration adjustment for 100 ms damping time (QFX5230-64CD, QFX5240-64OD, QFX5240-64QD)**—To achieve a damping time of 100 ms on a fully loaded setup with active hold timers on all ports, you must configure the hold timer to approximately 120 ms, accounting for an additional 20 ms. This adjustment ensures that the system meets the desired damping expectations effectively.

## Junos Telemetry Interface

- **Export packet capture statistics to external collectors using Junos Telemetry (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5700, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—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](#).

## Network Management and Monitoring

- **TAP aggregation (QFX5130-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Test access point (TAP) aggregation, similar to port mirroring, is a network monitoring and troubleshooting tool. Unlike port mirroring, TAP aggregation provides many-to-many packet replication, enabling you to capture different types of data in real time so that you quickly see what is happening in your network. You configure the TAP-aggregation feature at the `[edit forwarding-options tap-aggregation]` hierarchy level.

[See [TAP Aggregation for Network Monitoring](#).]

- **Dropped-packet notification (QFX5240-64OD and QFX5240-64QD)**—Packet drops are common occurrences on network switches and routers. Debugging packet drops can be complex and time-consuming. The packet-processing pipeline supports a limited set of drop counters, but these counters are insufficient for debugging complex packet-drop issues. Debugging difficulties can result in high mean times to recovery (MTTRs).

A feature called *dropped-packet notification*, also referred to as mirror on drop (MoD), can help you debug packet drops in real time. The areas of packet drop monitored include:

- Packets dropped due to processing in the ingress pipeline
- Packets dropped due to processing in the egress pipeline
- Packets dropped due to congestion in the MMU

Dropped-packet notification on the platforms named in this description is stateless and flow unaware.

You configure much of the dropped-packet notification feature at the `[edit forwarding-options mirror-profile]` hierarchy level.

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

You can 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. You must configure the PTP reference clock on the tap-aggregation switch and PTP must be running when the timestamp is inserted. Your tap aggregation switch must also sync the PTP FPGA's recovered time-of-day with the system chip's time-of-day. The command you use to enable the timestamp option is:

- 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)**—You can apply ingress ACL (access control list) UDF filtering on tap interfaces, which enables you to selectively choose specific traffic to be sent to the tool interfaces on a TAP-aggregation switch. If the ACL match and the TAP aggregation rule conflict, the ACL match takes precedence.

Also on these switches, you now 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
```

## Routing Protocols

- **BGP support for deterministic path forwarding (DPF) in a CLOS network (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, QFX5700E, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—BGP DPF partitions a physical fabric into multiple logical fabrics for distinct flows. Single hop External BGP (EBGP) might not meet all DC flow requirements, especially for drop and latency sensitive AI ML flows. BGP DPF colors the single-hop EBGP session on each link with a fabric color. If a link belongs to the red fabric, the EBGP session over the link is colored red. A route with no color community is advertised over any colored or uncolored EBGP sessions.

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

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

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

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

To enable GLB for multi-link on a 3-CLOS IP fabric, include the `glb-multilink-mode max-val/avg-val` statement at `[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 `[edit protocols bgp]` hierarchy level.

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

## Software Installation and Upgrade

- **Firmware upgrade support (QFX5241-32OD, QFX5241-64OD, and QFX5241-64QD)**—Use the `request system firmware upgrade` command to upgrade firmware.

[See [request system firmware upgrade \(Junos OS Evolved\)](#).]

- **Secure boot and common BIOS support (QFX5241-64OD and QFX5241-64QD)**—The secure boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. Secure boot is enabled by default on supported platforms.

[See [Junos OS Evolved Overview](#) and [request system firmware upgrade \(Junos OS Evolved\)](#).]

## Additional Features

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

- **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](#).]

- **Supported transceivers, optical interfaces, and DAC cables (QFX Series)**—Select your product in the [Hardware Compatibility Tool](#) (HCT) to view supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.
- **Support for 256-way ECMP** (QFX5130-32CD, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—Use this feature to increase the number of direct BGP peer connections, improve latency, and optimize data flow by configuring up to 256 ECMP next hops for external BGP peers.

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

## What's Changed

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

## Known Limitations

### IN THIS SECTION

- [EVPN | 29](#)
- [Interfaces and Chassis | 30](#)
- [Multicast | 30](#)
- [Optics | 30](#)

Learn about known limitations in Junos OS Evolved Release 23.4X100-D40 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.

## EVPN

- For VXLAN decapsulation at line rate, maximum pkt size is 1182Bytes, after which MMU drops might be observed at decap node. **PR1873759**
- On QFX5240, for packets undergoing VXLAN decapsulation, line-rate throughput is not achieved on 50G/100G ports for <=1518 bytes packet and 200G ports for <=2048 bytes packet. **PR1873579**

## Interfaces and Chassis

- When 40G and 100G xcvrs are used together on the 16x100G PIC of QFX5700 such that they are connected to the same vendor serial-deserializer core on the Packet Forwarding Engine side, then one or more of the links from these interfaces may not come up. With this PIC, ports 0-3, 4-7, 8-11 and 12-15 share one vendor serial-deserializer core each. **PR1867341**
- 2x200G channelization is not supported on the QFX5700 platform for the 23.4X100-D40-EVO release. **PR1884887**

## Multicast

- With IGMP Snooping in EVPN-VXLAN with Mrouter sync feature enabled with `clear bgp neighbor all` in router where the igmp query is learnt locally when this aggregated Ethernet goes from up to down state and all its mrouter states are removed and then the aggregated Ethernet comes up [after core isolation] and the traffic coming into the system from another esi-lag gets dropped till all the query is learnt in aggregated Ethernet (AE) and it becomes mrouter. This is the design limitation as it has core isolation in place and in this scenario we expect the traffic loss [12 sec] i.e till the time aggregated Ethernet (AE) has learnt all the incoming queries from all vlan and marks its ifl as mrouter interface in corresponding vlan. **PR1876740**

## Optics

- PRBS is not supported on the optics 'QSFP28-100G-LR'. This optics has PAM4 modulation on the Line side & NRZ modulation on the host side. As soon as prbs packets received from host side, module is not able to distribute the same on the line side due to modulation mismatch. Due to this, Signal Detect is brought down and the prbs packets are not sent end-to-end. **PR1894769**

## Open Issues

### IN THIS SECTION

- [General Routing | 31](#)



Learn about open issues in Junos OS Evolved Release 23.4X100-D40 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 Junos Evolved OS platforms QFX5K (QFX5230/QFX5240), ungraceful swapping a 400G DAC with a 400G DR4 optic (or vice-versa) might lead to traffic loss. **PR1862711**
- Invalid values are shown for Carrier transition counter sometimes on initiating `clear interfaces statistics all` command and `picd` app restart. This issue is observed on QFX5220-32CD, QFX5230-64CD, QFX5240-64OD. **PR1864375**

## 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 | 33](#)
- [Creating a Service Request with JTAC | 34](#)

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

22 December 2025—Revision 5, Junos OS Evolved Release 23.4X100-D40

21 September 2025—Revision 4, Junos OS Evolved Release 23.4X100-D40

18 August 2025—Revision 3, Junos OS Evolved Release 23.4X100-D40

6 August 2025—Revision 2, Junos OS Evolved Release 23.4X100-D40

31 July 2025—Revision 1, Junos OS Evolved Release 23.4X100-D40

---

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.