

# Release Notes

Published  
2024-04-30

## Junos OS Evolved Release 22.2X100-D20

---

### Introduction

Use these release notes to find new features, software limitations, and open issues for Junos OS Evolved Release 22.2X100-D20.

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

# Table of Contents

## Junos OS Evolved Release Notes for QFX5230 and QFX5240 Devices

### What's New | 1

#### What's New in 22.2x100-D20 | 1

Hardware | 2

Software | 10

Interfaces | 17

#### What's New in 22.2X100-D10 | 20

### What's Changed | 29

### Known Limitations | 29

### Open Issues | 29

### Licensing | 31

### Finding More Information | 32

### Requesting Technical Support | 33

### Revision History | 34

# Junos OS Evolved Release Notes for QFX5230 and QFX5240 Devices

## IN THIS SECTION

- [What's New | 1](#)
- [What's Changed | 29](#)
- [Known Limitations | 29](#)
- [Open Issues | 29](#)

These release notes accompany Junos OS Evolved Release 22.2X100-D20 for QFX5230 and QFX5240 switches. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

## What's New

## IN THIS SECTION

- [What's New in 22.2x100-D20 | 1](#)
- [What's New in 22.2X100-D10 | 20](#)

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

## What's New in 22.2x100-D20

## IN THIS SECTION

- [Hardware | 2](#)

●	Software   10
●	Interfaces   17

## Hardware

- **New QFX5240 switches (QFX Series)**—Starting in Junos OS Release 22.2X100-D20, we introduce the 800GbE data center switches, QFX5240-OD and QFX5240-QD. These switches offer 64 800GbE ports of OSFP and QSFP-DD. Using breakout cables, you can configure 64 ports of 800GbE, 128 ports of 400GbE, and 256 ports of 100GbE for QSFP5240-OD and QSFP5240-QD.

**Table 1: Features Supported on QFX5240**

Feature	Description
Chassis	<ul style="list-style-type: none"><li>• Support for inbuilt Routing Engine, Control Board, power supply, fan trays, FPCs, and PICs on QFX5240-64OD and QFX5240-64QD switches.</li></ul>

Table 1: Features Supported on QFX5240 *(Continued)*

Feature	Description
Class of service	<ul style="list-style-type: none"> <li>• Support for CoS features on Layer 2 and Layer 3 interfaces. Supported CoS features include: <ul style="list-style-type: none"> <li>• IPv4 and IPv6 unicast routing.</li> <li>• Classification and rewrite rules (DSCP, IEEE-802.1p)</li> <li>• Port scheduling</li> <li>• Shared buffer</li> <li>• Priority-based Flow Control (PFC) based on IEEE-802.1p. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2).</li> <li>• Weighted Random Early Drop (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> </li> </ul> <p>[See <a href="#">Understanding How Class of Service Manages Congestion</a>.]</p>
Forwarding and sampling	<ul style="list-style-type: none"> <li>• Support for dynamic load balancing (DLB) (for port speeds over 50G) and resilient hashing (RH) for equal-cost multipath (ECMP) routes. DLB and RH are not supported on Link Aggregation Group (LAG), or when an LAG is one of the egress ECMP members.</li> </ul> <p>[See <a href="#">Dynamic Load Balancing</a>, <a href="#">Use of Resilient Hashing to Minimize Flow Remapping</a>, and <a href="#">ecmp-resilient-hash</a>.]</p>

Table 1: Features Supported on QFX5240 (Continued)

Feature	Description
Interfaces	<ul style="list-style-type: none"> <li>• <b>Interfaces support</b>—The QFX5240 switches have 64x800GbE OSFP ports on QFX5240-64OD, and 64x800GbE QSFP-DD ports on QFX5240-64QD. The last two ports (64 and 65) are 2x10GbE SFP28 on both the QFX5240 variants.</li> </ul> <p>With break-out cables, the ports support the following configurations:</p> <ul style="list-style-type: none"> <li>• 64x800GbE</li> <li>• 128x400GbE</li> <li>• 256x100GbE</li> </ul> <p>QFX5240 also supports 2x400GbE and 8x100GbE channelizations.</p> <p><b>NOTE:</b> For 8x100GbE channelization, only even ports can be channelized with a pair port marked as unused. (using set groups global interfaces et-0/0/&lt;port&gt; unused command). Refer <a href="#">"List item." on page 17</a> for a mapping of pair ports on the OSFP and QSPDD-800 ports.</p> <p>Refer the <a href="#">Port Checker tool</a> for information on different port combinations supported on the QFX5240 switches.</p> <p><b>NOTE:</b> On the QFX5240 switches, the runts (under Input errors) and fragment frames (under MAC statistics) counters do not increment in the output of the <b>show interfaces extensive</b> command. These counters are not supported due to a chip limitation.</p> <p>[See <a href="#">Port Settings.</a>]</p>

**Table 1: Features Supported on QFX5240 (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. This includes:<ul style="list-style-type: none"><li>• IGMP snooping with 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) IGMP snooping</li><li>• IGMP snooping with integrated routing and bridging (IRB)</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: Features Supported on QFX5240 (Continued)

Feature	Description
Layer 3 features	<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> <li>• Support for Layer 3 unicast forwarding and generic routing encapsulation (GRE) tunneling. Both IPv4 and IPv6 unicast routing are supported. [See <a href="#">Generic Routing Encapsulation (GRE)</a>.]</li> <li>• Support for Layer 3 multicast forwarding. This includes: <ul style="list-style-type: none"> <li>• PIM first hop router (FHR), Rendezvous point (RP) functionality</li> <li>• Multicast Source Discovery Protocol (MSDP)</li> <li>• PIM Dense Mode (PIM DM)</li> <li>• Make before break (MBB) support for multicast receivers on existing Layer 3 aggregated Ethernet (aex) or link aggregation group (LAG) interfaces, which includes member addition, deletion, link up, and down events.</li> <li>• Protocol Independent Multicast - source-specific multicast (PIM SSM)</li> <li>• Protocol Independent Multicast - sparse mode (PIM SM)</li> <li>• L3 multicast forwarding on integrated routing and bridging (IRB) interfaces: <ul style="list-style-type: none"> <li>• IPv4 and IPv6 multicast</li> <li>• Internet Group Management Protocol (IGMP v1/v2/v3)</li> <li>• Multicast Listener Discovery (MLD v1/v2)</li> <li>• Any-source multicast (ASM) and source-specific multicast (SSM) modes</li> </ul> </li> </ul> </li> </ul> <p>[See <a href="#">Multicast Routing Protocols</a> and <a href="#">PIM Overview</a>.]</p>



Table 1: Features Supported on QFX5240 (Continued)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> <li>Support for sFlow. [See <a href="#">Overview of sFlow Technology</a>.]</li> <li>Support for analyzers and port mirroring. 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 for configuring firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See <a href="#">Overview of JET APIs</a>.]</li> <li>Platform resiliency support on QFX5240-64OD and QFX5240-64QD switches for hardware components of each FRU. If a failure is detected on a hardware component, the software: <ol style="list-style-type: none"> <li>1. Logs the message to give clear indication of failure details, including time stamp, module name, component name and failure details.</li> <li>2. Raises SNMP trap.</li> <li>3. Glows the FRU fault if LED is present.</li> <li>4. Performs local action such as self-healing, or taking the component out of service.</li> </ol> </li> </ul>
Protection against DDoS attacks	<ul style="list-style-type: none"> <li>Supports configuration and installation of policers at the Packet Forwarding Engine (PFE) level for defense from DDoS attacks. By default, DDoS protection is enabled for many protocols on the QFX5240-64OD and QFX5240-64QD 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>

Table 1: Features Supported on QFX5240 *(Continued)*

Feature	Description
Routing policy and firewall filters	<ul style="list-style-type: none"> <li>• Firewall filter support on Layer 2 and Layer 3 interfaces.</li> </ul> <p>[See <a href="#">Firewall Filter Match Conditions and Actions</a> and <a href="#">Configuring Enhanced Egress Firewall Filters</a>.]</p>
Services applications	<ul style="list-style-type: none"> <li>• Support for generic routing encapsulation (GRE) features: <ul style="list-style-type: none"> <li>• GRE tunnels over GigE, LAG, and VLAN</li> <li>• Tagged sub-interfaces</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>

Table 1: Features Supported on QFX5240 (Continued)

Feature	Description
Software installation and upgrade	<ul style="list-style-type: none"> <li>Firmware upgrade support. The following commands are supported: <ul style="list-style-type: none"> <li><code>request system firmware upgrade re bios</code></li> <li><code>request system firmware upgrade re fpga</code></li> <li><code>request system firmware upgrade re gfpga (QFX5240-64OD FPGA)</code></li> <li><code>request system firmware upgrade re ssd</code></li> <li><code>request system firmware upgrade fpc opticscpld&lt;0/1/2&gt;</code></li> <li><code>request system firmware upgrade re fancpld</code></li> <li><code>request system firmware upgrade re i210</code></li> <li><code>request system firmware upgrade fpc bcm-pfe</code></li> </ul> <p>[See <a href="#">request system firmware upgrade</a>.]</p> </li> <li>Support for zero-touch provisioning (ZTP) over IPv4 and IPv6 on the management and WAN interfaces. <p>[See <a href="#">Zero Touch Provisioning</a>.]</p> </li> <li>Support for USB booting.</li> </ul> <p><b>NOTE:</b> On QFX5240 switches, only UEFI boot media (UEFI USB, UEFI NVME, UEFI Network, and so on) is supported. You must select USB (UEFI USB) manually from the BIOS menu, or use the <b><code>request node reboot re0 usb</code></b> command to boot from USB.</p>
Support for optics	<ul style="list-style-type: none"> <li><b>Supported transceivers, optical interfaces, and DAC cables</b>—Select your product in the <a href="#">Hardware Compatibility Tool</a> to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.</li> </ul>

## Software

### IN THIS SECTION

- [Additional features optimized for AI ML fabrics | 10](#)

### Additional features optimized for AI ML fabrics

- **Reactive Path Rebalancing** (QFX5240)—Starting in 22.2X100-D20 release, QFX5240 devices support Reactive Path Rebalancing. Reactive Path Rebalancing is an enhancement to the existing Flowlet mode in the Dynamic Load Balancing (DLB) feature. In the Flowlet mode of DLB, the user configures an inactivity interval. The traffic uses assigned outgoing interface until a pause in flow is greater than the inactivity timer. It is possible that the current outgoing link quality becomes worse over a period of time and the pause within the flow does not exceed the inactivity timer that is configured. Classic Flowlet mode does not reassign to a different link within the inactivity interval and cannot utilize a better quality link. Reactive path rebalancing addresses this limitation by enabling the user to move the traffic to a link with a better quality in the Flowlet mode.

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

Benefits of the reactive path load balancing are:

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

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


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

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

## Example

Consider topology as shown in [Figure 1 on page 11](#), where there are three ingress ports and two egress ports in a device. Also shown are table entries forwarding the traffic to each of the egress ports. All the ingress and egress ports are of the same speed.

**Figure 1: Reactive Path Rebalancing**



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

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

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

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

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

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

- **PFC watchdog support (QFX5230-64CD, QFX5240-64OD, QFX5240-64QD)** —Starting with Junos OS Evolved Release 22.2X100-D20, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD switches support the PFC watchdog feature. The PFC watchdog monitors PFC-enabled ports for PFC pause storms. When a PFC-enabled port receives PFC pause frames for an extended period of time and PFC watchdog does not detect flow control frames on that port, PFC watchdog mitigates the situation. It does this by disabling the queue where the PFC pause storm was detected for a configurable length of time called the recovery time. After the recovery time passes, PFC watchdog re-enables the affected queue.

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

- **poll-interval**—The interval at which PFC watchdog checks the status of PFC queues, which can be 1, 10, or 100 milliseconds.
- **detection**—The number of polling intervals the PFC watchdog waits before it mitigates the stalled traffic, from 1-15 intervals.
- **watchdog-action**—The action the PFC watchdog takes to mitigate a stalled traffic queue, either drop or forward all enqueued and newly arriving packets.
- **recovery**—How long the PFC watchdog disables the affected queue before it restores PFC on the queue, from 100-1500 milliseconds with a default of 200 milliseconds.

[See [PFC Watchdog](#) and [congestion-notification-profile](#).]

- **Priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged IPv6 traffic (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—To support lossless IPv6 traffic across Layer 3 connections to Layer 2 subnetworks, you can configure PFC to operate using 6-bit DSCP values from Layer 3 headers of untagged VLAN traffic, rather than IEEE 802.1p priority values in Layer 2 VLAN-tagged packet headers. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2). To enable DSCP-based PFC, map a forwarding class to a PFC priority using the `pfc-priority` statement, define a congestion notification profile to enable PFC on traffic specified by a 6-bit DSCP value, and set up a classifier for the DSCP value and the PFC-mapped forwarding class.

[See [Understanding PFC Using DSCP at Layer 3 for Untagged Traffic](#).]

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

[See [xon \(Input Congestion Notification\)](#).]

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

[See [buffer-dynamic-threshold](#).]

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

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

[See [Configuring Ingress and Egress Dedicated Buffers](#).]

- **egress-quantization (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)** - Starting in Junos EVO 22.2X100-D20, you can modify port load and port queue metrics from their default values so that when dynamic load balancing is enabled, the metrics are used to determine an optimal link. Use the new `egress-quantization` CLI to configure the desired ratio of port load metric to port queue metric based on the traffic pattern.

[See [egress-quantization](#).]

- **Enhanced telemetry legacy gRPC dial-out IP address and routing instance support (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Junos OS Evolved Release 22.2X100-D20 enhances legacy telemetry dial-out to support outgoing gRPC requests from a device's specified IP address or interface address (such as loopback 0). Currently, the outgoing interface IP address is used as the source address. To instead configure the source IP address, include the statement `export-profile ep1 local-address ip-address` at the `[edit services analytics ]` hierarchy level. IPv6 addresses are supported. Configuring an IPv6 address is only supported for legacy gRPC dial-out connections. If configured for a UDP transport, an error is returned. You also can configure the routing instance for gRPC dial-out connections. This is optional. To set the routing instance name, include the statement `export-profile ep1 routing-instance routing-instance-name` at the `[edit services analytics ]` hierarchy level. If the routing instance is not configured, the default routing instance is applied.

[For sensors, see [Junos YANG Data Model Explorer](#). To configure the export profile, see [export-profile \(Junos Telemetry Interface\)](#).]

- **Telemetry for IPv4 and IPv6 traffic statistics (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Junos OS Evolved Release 22.2X100-D20 introduces streaming support for IPv4 and IPv6 transit statistics using the native resource path `/junos/system/linecard/interface/traffic` or the OpenConfig resource path `/interfaces/interface/`. This feature supports rate counters that can help you diagnose real time issues. The following fields are exported:

- `if_in_ipv4pkts`
- `if_in_ipv4_1sec_pkts`

- if\_in\_ipv4\_bytes
- if\_in\_ipv4\_1sec\_octets
- if\_out\_ipv4pkts
- if\_out\_ipv4\_1sec\_pkts
- if\_out\_ipv4\_bytes
- if\_out\_ipv4\_1sec\_octets
- if\_in\_ipv6pkts
- if\_in\_ipv6\_1sec\_pkts
- if\_in\_ipv6\_bytes
- if\_in\_ipv6\_1sec\_octets
- if\_out\_ipv6pkts
- if\_out\_ipv6\_1sec\_pkts
- if\_out\_ipv6\_bytes
- if\_out\_ipv6\_1sec\_octets

To enable transit statistics for the physical port you must configure route accounting. To enable IPv4 route accounting, include the `inet4 route-accounting` statement at the `[edit forwarding-options family ]` hierarchy level. To enable IPv6 route accounting, include the `inet6 route-accounting` statement at the `[edit forwarding-options family ]` hierarchy level.

[For sensors, see [Junos YANG Data Model Explorer](#). To configure route accounting, see [route-accounting](#). For Junos Telemetry interface configuration, see [Junos Telemetry Interface User Guide](#).]

- **Extended sFlow Functionality Support** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—Starting in Junos OS Evolved Release 22.2X100-D20, we've extended the sflow monitoring functionality to support the following features:

- Export of sFlow sample packets via `mgmt_junos` interface.

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



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

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

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

sFlow is a traffic monitoring protocol that supports VRFs. sFlow provides traffic sampling on configured ports based on sample rate and port information to a collector. An sFlow monitoring system consists of an sFlow agent embedded in the device and up to four external collectors. The sFlow agent performs packet sampling and gathers interface statistics, and then combines the information into UDP datagrams that are sent to the sFlow collectors.

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

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

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

See [collector](#) and [show sflow collector](#).

- **Remote port mirroring to IPv4/IPv6 address (GRE encapsulation) with DSCP, source-address, and rate-limiting parameters (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)**—Starting in Junos OS Evolved Release 22.2X100-D20, you can configure DSCP, source-address, and rate-limiting parameters in your configuration for remote port mirroring to IPv4 or IPv6 addresses. You use remote port mirroring to copy packets entering a port or VLAN and send the copies to the IPv4 or IPv6 address of a device running an analyzer application on a remote network (sometimes referred to as "extended remote port mirroring"). The mirrored packets are GRE-encapsulated.

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

```
[edit forwarding-options analyzer instance instance-name output]
```

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

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

```
set class-of-service forwarding-classes class class-name queue-num queue-number
```

```
set class-of-service interfaces interface-name scheduler-map map-name
```

```
set class-of-service scheduler-maps map-name forwarding-class class-name scheduler scheduler-name
```

```
set class-of-service schedulers scheduler-name shaping-rate rate
```

[See [Port Mirroring and Analyzers](#).]

- **SNMP support for PFC, ECN, and CoS ingress packet drop accounting** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—Junos OS Evolved Release 22.2X100-D20 introduces SNMP support that helps to account for the packets that are dropped because of ingress port congestion. You can view and export the error counters data for explicit congestion notification (ECN), ingress drops, and priority-based flow control (PFC) using the following commands:

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

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

- **Telemetry support for PFC, ECN, and CoS ingress packet drop accounting** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—Junos OS Evolved Release 22.2X100-D20 introduces telemetry support to account for packets that are dropped due to ingress port congestion. Counters for priority flow control (PFC), explicit congestion notification (ECN), and ingress drops are exported using the sensor `/junos/system/linecard/interface/traffic`.

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

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

- **rdma-opcode firewall filter match condition** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD) - Starting in Junos EVO 22.2X100-D20, `rdma-opcode` and `rdma-opcode-except` firewall filter match conditions have been added to enable match on InfiniBand Base Transport header opcode.

[See [rdma-opcode](#).]

- **Selectively enable or disable dynamic load balancing using `dynamic-load-balance`** (QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD) - Starting in Junos EVO 22.2X100-D20, you can selectively enable or disable dynamic load balancing based on any match condition available in firewall filters using the new `dynamic-load-balance` CLI.

[See [dynamic-load-balance](#).]

- **Strip and replace BGP private-AS path** (QFX5230-64CD, QFX5240-64OD, QFX5240-64QD)—Starting in Junos OS Evolved Release 22.2X100-D20, we have introduced the `strip-as-path` policy

option that removes the incoming autonomous system (AS) path as part of the import policy for a BGP session and replaces the received autonomous system (AS) path with the receiving router's local AS number for the receiving session. Note that the local AS number may be different from the number configured under `autonomous system` in the `[edit routing-options]` hierarchy.

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

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

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

[See [Autonomous Systems for BGP Sessions](#).]

## Interfaces

- **Channelization Support (QFX5240-64OD, QFX5240-64QD)**—Starting in Junos EVO 22.2X100-D20, the QFX5240-64OD and QFX5240-64QD switches support 2x400GbE and 8x100GbE channelizations.

For the 8x channelized mode on the OSFP and QSFPDD-800 ports, only even ports can be channelized. When an even port is channelized, a pair port must be marked as unused.

[Table 2 on page 17](#) displays the mapping of pair ports on the OSFP and QSFPDD-800 ports.

**Table 2: Channelization Port Mapping**

Ports for 8x Channelization	Pair port to mark unused on OSFP	Pair port to mark unused on QSFPDD-800
et-0/0/0	et-0/0/33	et-0/0/1
et-0/0/2	et-0/0/35	et-0/0/3
et-0/0/4	et-0/0/37	et-0/0/5
et-0/0/6	et-0/0/39	et-0/0/7
et-0/0/8	et-0/0/41	et-0/0/9

**Table 2: Channelization Port Mapping (Continued)**

Ports for 8x Channelization	Pair port to mark unused on OSFP	Pair port to mark unused on QSFPDD-800
et-0/0/10	et-0/0/43	et-0/0/11
et-0/0/12	et-0/0/45	et-0/0/13
et-0/0/14	et-0/0/47	et-0/0/15
et-0/0/16	et-0/0/49	et-0/0/17
et-0/0/18	et-0/0/51	et-0/0/19
et-0/0/20	et-0/0/53	et-0/0/21
et-0/0/22	et-0/0/55	et-0/0/23
et-0/0/24	et-0/0/57	et-0/0/25
et-0/0/26	et-0/0/59	et-0/0/27
et-0/0/28	et-0/0/61	et-0/0/29
et-0/0/30	et-0/0/63	et-0/0/31
et-0/0/32	et-0/0/1	et-0/0/33
et-0/0/34	et-0/0/3	et-0/0/35
et-0/0/36	et-0/0/5	et-0/0/37
et-0/0/38	et-0/0/7	et-0/0/39

**Table 2: Channelization Port Mapping (Continued)**

Ports for 8x Channelization	Pair port to mark unused on OSFP	Pair port to mark unused on QSPDD-800
et-0/0/40	et-0/0/9	et-0/0/41
et-0/0/42	et-0/0/11	et-0/0/43
et-0/0/44	et-0/0/13	et-0/0/45
et-0/0/46	et-0/0/15	et-0/0/47
et-0/0/48	et-0/0/17	et-0/0/49
et-0/0/50	et-0/0/19	et-0/0/51
et-0/0/52	et-0/0/21	et-0/0/53
et-0/0/54	et-0/0/23	et-0/0/55
et-0/0/56	et-0/0/25	et-0/0/57
et-0/0/58	et-0/0/27	et-0/0/59
et-0/0/60	et-0/0/29	et-0/0/61
et-0/0/62	et-0/0/31	et-0/0/63

Refer the [Port Checker tool](#) for information on different port combinations supported on the QFX5240 switches.

[See [Port Settings](#).]

# What's New in 22.2X100-D10

IN THIS SECTION

- [Hardware | 20](#)

## Hardware

- **New QFX5230-64CD switch (QFX Series)**—Starting in Junos OS Evolved Release 22.2X100-D10, QFX5230-64CD offers high-density 400-GbE access ports in a QSFP-DD form factor optimized for high-end spine and super-spine layer of the IP fabric multitier architecture in a 2-RU fixed form factor. The QFX5230-64CD switch provides a uni-directional throughput of 25.6 Tbps and offers sixty four 400-GbE network ports and up to 128x200GbE, 256x100GbE, 64x40GbE, 256x25GbE, and 258x10GbE ports.

Table 3: Features Supported on QFX5230-64CD

Feature	Description
Chassis	<ul style="list-style-type: none"><li>• Support for inbuilt Routing Engine, Control Board, power supply, fan trays, FPCs, and PICs on QFX5230-64CD devices.</li></ul>

Table 3: Features Supported on QFX5230-64CD (Continued)

Feature	Description
Class of service	<ul style="list-style-type: none"> <li>• Support for CoS features on Layer 2 and Layer 3 interfaces. Both IPv4 and IPv6 unicast routing are supported. Other supported CoS features include: <ul style="list-style-type: none"> <li>• Classification and rewrite rules (DSCP, IEEE-802.1p)</li> <li>• Port scheduling</li> <li>• Shared buffer</li> <li>• Priority-based Flow Control (PFC) based on IEEE-802.1p. DSCP-based PFC is required to support Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2).</li> <li>• Weighted Random Early Drop (WRED) and Explicit Congestion Notification (ECN)</li> <li>• Telemetry support for CoS queue statistics exported using the sensor /junos/system/linecard/qmon-sw/.</li> </ul> </li> </ul> <p>[See <a href="#">Understanding How Class of Service Manages Congestion</a>.]</p>
Forwarding and sampling	<ul style="list-style-type: none"> <li>• Support for dynamic load balancing (DLB) and resilient hashing (RH) for equal-cost multipath (ECMP) routes. DLB and RH are not supported on Link Aggregation Group (LAG).</li> </ul> <p>[See <a href="#">Dynamic Load Balancing</a>, <a href="#">Use of Resilient Hashing to Minimize Flow Remapping</a>, and <a href="#">ecmp-resilient-hash</a>.]</p>

Table 3: Features Supported on QFX5230-64CD (Continued)

Feature	Description
Interfaces	<ul style="list-style-type: none"> <li>• <b>Interfaces Support</b>—QFX5230-64CD has 64 QSFP56-DD ports and two SFP+ ports. The QSFP56-DD ports support the following speeds: <ul style="list-style-type: none"> <li>• 400 GbE</li> <li>• 200 GbE</li> <li>• 100 GbE</li> <li>• 50GbE</li> <li>• 40 GbE</li> </ul> </li> </ul> <p>The QSFP-DD ports also support the following speeds (with breakout cables):</p> <ul style="list-style-type: none"> <li>• 50GbE</li> <li>• 25GbE</li> <li>• 10GbE</li> </ul> <p>The SFP+ ports support 10 GbE.</p> <p>QFX5230-64CD supports 1x400GbE, 2x200GbE, 4x100GbE, 2x100GbE, 1x100GbE, 2x50GbE, 1x50GbE, 1x40GbE, 4x25GbE, and 4x10GbE channelizations.</p> <p>[See <a href="#">Port Settings</a>.]</p>



Table 3: Features Supported on QFX5230-64CD (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. This includes: <ul style="list-style-type: none"> <li>• IGMP snooping with 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) IGMP snooping</li> <li>• IGMP snooping with integrated routing and bridging (IRB)</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 3: Features Supported on QFX5230-64CD (Continued)

Feature	Description
Layer 3 features	<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> <li>• Support for Layer 3 unicast forwarding and generic routing encapsulation (GRE) tunneling. Both IPv4 and IPv6 unicast routing are supported.  [See <a href="#">Generic Routing Encapsulation (GRE)</a>.]</li> <li>• Support for Layer 3 multicast forwarding. This includes: <ul style="list-style-type: none"> <li>• PIM first hop router (FHR), Rendezvous point (RP) functionality</li> <li>• Multicast Source Discovery Protocol (MSDP)</li> <li>• PIM Dense Mode (PIM DM)</li> <li>• Make before break (MBB) support for multicast receivers on existing Layer 3 aggregated Ethernet (aex) or link aggregation group (LAG) interfaces, which includes member addition, deletion, link up, and down events</li> <li>• Protocol Independent Multicast - source-specific multicast (PIM SSM)</li> <li>• Protocol Independent Multicast - sparse mode (PIM SM)</li> <li>• L3 multicast forwarding on integrated routing and bridging (IRB) interfaces: <ul style="list-style-type: none"> <li>• IPv4 and IPv6 multicast</li> </ul> </li> </ul> </li> </ul>

Table 3: Features Supported on QFX5230-64CD *(Continued)*

Feature	Description
	<ul style="list-style-type: none"> <li>• Internet Group Management Protocol (IGMP v1/v2/v3)</li> <li>• Multicast Listener Discovery (MLD v1/v2)</li> <li>• Any-source multicast (ASM) and source-specific multicast (SSM) modes</li> </ul> <p>[See <a href="#">Multicast Routing Protocols</a> and <a href="#">PIM Overview</a>.]</p>
Network management and monitoring	<ul style="list-style-type: none"> <li>• Support for sFlow.</li> </ul> <p>[See <a href="#">Overview of sFlow Technology</a>.]</p> <ul style="list-style-type: none"> <li>• Support for Analyzers and Port Mirroring.</li> </ul> <p>[See <a href="#">Understanding Port Mirroring and Analyzers</a>.]</p>

Table 3: Features Supported on QFX5230-64CD (Continued)

Feature	Description
Platform and infrastructure	<ul style="list-style-type: none"> <li>Platform resiliency support for QFX5230-64CD switches, for hardware components of each FRU. If a failure is detected on a hardware component, the software: <ul style="list-style-type: none"> <li>Logs the message to give clear indication of failure details, including time stamp, module name, component name &amp; failure details.</li> <li>Raises/clears alarms, if applicable.</li> <li>Raises SNMP trap.</li> <li>Glow the FRU fault if LED is present.</li> <li>Performs local action such as self-healing or taking the component out of service.</li> </ul> </li> <li>Support for configuring firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs.</li> </ul> <p>[See <a href="#">Overview of JET APIs</a>.]</p>
Protection against DDoS attacks	<ul style="list-style-type: none"> <li>Supports configuration and installation of policers at the Packet Forwarding Engine (PFE) level for defense from DDoS attacks. By default, DDoS protection is enabled for many protocols on the QFX5230-64CD switches.</li> </ul> <p>[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>.]</p>
Routing policy and firewall filters	<ul style="list-style-type: none"> <li>Firewall filter support on Layer 2 and Layer 3 interfaces.</li> </ul> <p>[See <a href="#">Firewall Filter Match Conditions and Actions</a> and <a href="#">Configuring Enhanced Egress Firewall Filters</a>.]</p>

**Table 3: Features Supported on QFX5230-64CD (Continued)**

Feature	Description
Services applications	<ul style="list-style-type: none"><li>• Support for generic routing encapsulation (GRE) features:<ul style="list-style-type: none"><li>• GRE tunnels over GigE, LAG, and VLAN</li><li>• Tagged sub-interfaces</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>

Table 3: Features Supported on QFX5230-64CD (Continued)

Feature	Description
Software Installation and Upgrade	<ul style="list-style-type: none"> <li>Firmware upgrade support. The following commands are supported: <ul style="list-style-type: none"> <li>request system firmware upgrade re bios</li> <li>request system firmware upgrade re fpga</li> <li>request system firmware upgrade re i210</li> <li>request system firmware upgrade re ssd</li> <li>request system firmware upgrade re xmcfgpa</li> <li>request system firmware upgrade re fancpld</li> <li>request system firmware upgrade fpc slot 0 bcm-pfe</li> <li>request system firmware upgrade fpc slot 0 opticscpld&lt;0 1 2&gt;</li> </ul> <p>[See <a href="#">request system firmware upgrade</a>.]</p> </li> <li>Support for secure BIOS and secure boot implementation based on the UEFI 2.4 standard. <p>[See <a href="#">Secure Boot</a>.]</p> </li> <li>Zero Touch Provisioning support for WAN interfaces and DHCPv6 options. <p>[See <a href="#">Zero Touch Provisioning</a>.]</p> </li> </ul>
Support for optics	<ul style="list-style-type: none"> <li><b>Supported transceivers, optical interfaces, and DAC cables</b>—Select your product in the <a href="#">Hardware Compatibility Tool</a> to view supported transceivers, optical interfaces, and DAC cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.</li> </ul>

## 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 22.2X100-D20 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.

## Open Issues

### IN THIS SECTION

- [Class of Service | 29](#)
- [General Routing | 30](#)
- [Interfaces and Chassis | 31](#)
- [Network Management and Monitoring | 31](#)

Learn about open issues in Junos OS Evolved Release 22.2X100-D20 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.

## Class of Service

- On 22.2X100-D20 release, when cos BA classifier entries are modified (specifically entry delete) on the fly along with forwarding class modification through `load override <>` option, cosd process will core and will fail to start until device reboot performed. This issue will not be seen with regular CLI configuration flow or with configuration rollback. Workaround to avoid this problem is, while

performing load override <> if there is any existing classifier entries are deleted from configuration, deactivate that classifier along with bind before the configuration commit. Activate the classifier and bind after the commit. By this way classifier configuration modification segregated from load override <>. [PR1805461](#)

## General Routing

- Configuring vlan id 1 as below for Layer 3 subinterface configuration as below might not work as expected on 22.2X100 branch. set interfaces et-0/0/64 unit 0 vlan-id 1 <<<< unsupported config vlan id 1 set interfaces et-0/0/64 unit 0 family inet address 20.0.0.2/30 set interfaces et-0/0/64 unit 0 family iso set interfaces et-0/0/64 unit 0 family inet6 address 2254:254:13::2/126. Customer can use vlan id other than 1. [PR1765945](#)
- QFX5240-64OD and QFX5240-64QD: Any port configuration changes (such as channelisation, port-speed) of the peer port when traffic is running causes few CRC errors on adjacent ports. Recommendation is to clear the statistics on the ports to reset the CRC error counters after any port configuration change. [PR1770589](#)
- Back-to-back speed or channelization configuration changes within a short interval (less than 3 minutes) on same interface can cause optic to get stuck during initialization. [PR1775128](#)
- On Junos OS Evolved QFX platforms QFX5230-64CD, QFX5240-64OD and QFX5240-64QD, inserting and removing the USB flash device and running show hardware chassis details command will lead to show the details of USB even after removing it from the device. [PR1793934](#)
- On Junos OS Evolved QFX platforms, when CPU board EEPROM device fails, there is no alarm or logs seen indicating failure of hardware. [PR1791331](#)
- Application might restart due to this issue after reboot. [PR1794769](#)
- IPv6 Traffic Statistics show large/invalid values at times when route-accounting is deactivated and activated. Steps that lead to this issue:
  1. Configure forwarding-options route-accounting inet6.
  2. Send ipv6 traffic. 3. Clear interface statistics
  3. Delete forwarding-options route-accounting inet6
  4. Add forwarding-options route-accounting inet6
  5. IPv6 CLI statistics are now showing large/invalid values

If IPv6 statistics under show interfaces show large or invalid values, workaround is to execute clear interfaces statistics all after the issue has occurred. [PR1798636](#)



- Filters with multiple UDF matches and UDF bit-length  $\geq 16$ , might not work after an update to the UDF match due to BCM SDK issue with UDF group configuration update. [PR1797665](#)

## Interfaces and Chassis

- On all Junos OS Evolved platforms, if changing of the speed and creation of the aggregated Ethernet interfaces is in the same commit, the commit fails with **Interface aeX with child links of mixed speed but link-speed mixed is not configured**. [PR1743461](#)
- On QFX5230-64CD 400G DAC cable of 2.5m and 4X100G DAC, break out might not link up with some peer devices. This issue is not seen with all peer devices, Recommendation for this release is use 1m DAC cable or supported 400G Optics. [PR1747315](#)
- On IRB interface, logical interface level statistics not supported on QFX5230 and QFX5240 platforms due to software limitation. [PR1772350](#)
- With sflow egress sampling on VLAN tagged/aggregated Ethernet interface, does not have the correct VLAN info in the samples. [PR1775904](#)
- On QFX5000 Junos OS Evolved platforms, when a USB is plugged into the chassis, the USB disk does not show up in show chassis hardware detail CLI command output until a subsequent reboot. Similarly when the USB stick is plugged out of the chassis, the show chassis hardware detail CLI command continues to show up the USB stick as plugged into the chassis until subsequent reboot of the device. [PR1793934](#)
- Back-to-back speed or channelization configuration changes within a short interval (less than 3 minutes) on same interface can cause optic to get stuck during initialization. [PR1797665](#)

## Network Management and Monitoring

Sample size output is not displayed in show sflow output. [PR1772141](#)

## 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

1 May 2024—Revision 1, Junos OS Evolved Release 22.2X100-D20

---

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 © 2024 Juniper Networks, Inc. All rights reserved.