

Release Notes

Published
2026-09-29

Junos OS Evolved 25.2X100-D20

Table of Contents

Introduction | 1

Junos OS Evolved Release Notes for QFX Series

What's New in 25.2X100-D20 | 1

- Hardware | **2**
- Chassis | **18**
- Class of Service (CoS) | **19**
- EVPN | **20**
- Forwarding Options | **22**
- High Availability (HA) and Resiliency | **22**
- Interfaces | **22**
- Network Management and Monitoring | **23**
- Precision Time Protocol (PTP) | **25**
- Routing Protocols | **25**
- Services Applications | **26**
- Additional Features | **26**

What's Changed in 25.2X100-D20 | 29

- Services Applications | **30**

Known Limitations in 25.2X100-D20 | 30

- General Routing | **30**

Open Issues in 25.2X100-D20 | 31

- General Routing | **31**
- Interfaces and Chassis | **32**

Migration, Upgrade, and Downgrade Instructions | 33

Documentation Updates | 33

Licensing | 34

Finding More Information | 34

Requesting Technical Support | 35

Revision History | 37

Introduction

Junos OS Evolved runs on the following Juniper Networks® hardware: ACX Series, PTX Series and QFX Series. Use these release notes to find new and updated features, software limitations, and open issues in Junos OS Evolved Release 25.2X100-D20 for QFX Series.

You can find release notes for all Junos OS Evolved releases at [Junos OS Evolved Documentation | Juniper Networks](#)

Junos OS Evolved Release Notes for QFX Series

IN THIS SECTION

- [What's New in 25.2X100-D20 | 1](#)
- [What's Changed in 25.2X100-D20 | 29](#)
- [Known Limitations in 25.2X100-D20 | 30](#)
- [Open Issues in 25.2X100-D20 | 31](#)

What's New in 25.2X100-D20

IN THIS SECTION

- [Hardware | 2](#)
- [Chassis | 18](#)
- [Class of Service \(CoS\) | 19](#)
- [EVPN | 20](#)
- [Forwarding Options | 22](#)
- [High Availability \(HA\) and Resiliency | 22](#)
- [Interfaces | 22](#)

- Network Management and Monitoring | **23**
- Precision Time Protocol (PTP) | **25**
- Routing Protocols | **25**
- Services Applications | **26**
- Additional Features | **26**

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

Hardware

New QFX Series Switch (QFX5140-24CD8O)—The Juniper Networks® QFX5140-24CD8O switch is a next-generation 1RU data center switch designed for high-density leaf and spine deployments. It delivers 16 Tbps of line-rate throughput, supporting 24x400G QSFP112 and 8x800G OSFP ports with flexible breakout and mixed-speed operation. QFX5140-24CD8O enables scalable, future-ready data center fabrics with support for advanced optics, low-latency performance, and adaptable airflow and power options, including ORV3 open-rack configurations.

Table 1:QFX5140-24CD8O Feature Support

Feature	Description
Bidirectional Forwarding Detection	• Support for BFD (OSPF, IS-IS, and BGP). [See Understanding BFD.]

(Continued)

Feature	Description
Class of service (CoS)	• Support for CoS features on Layer 2 (L2) and Layer 3 (L3) interfaces, including: • IPv4 and IPv6 unicast routing • Classification and rewrite rules for Differentiated Services code point (DSCP) and IEEE-802.1p • Port scheduling • Shared buffer • Priority-based flow control (PFC) based on IEEE-802.1p for VLAN-tagged traffic. Protocols such as Remote Direct Memory Access over Converged Ethernet version 2 (RoCEv2) require DSCP-based PFC at L3 for untagged traffic. • Weighted random early detection (WRED) and explicit congestion notification (ECN) • Telemetry support for CoS queue statistics exported using the sensor / <code>junos/system/linecard/qmon-sw/</code>. [See Traffic Management User Guide (QFX Series Switches and EX4600 Switches).]
Dynamic Host Configuration Protocol (DHCP)	• Support for DHCPv4 and DHCPv6 stateless relay. [See DHCP Relay Agent.]

(Continued)

Feature	Description
EVPN	• Support for the following features: • CoS, and firewall filtering and policing on EVPN-VXLAN [See CoS Support on EVPN VXLANs and Firewall Filter Match Conditions and Actions (QFX and EX Series Switches).] • sFlow technology on EVPN-VXLAN [See Overview of sFlow Technology.] • Port mirroring and analyzers on EVPN-VXLAN. [See Port Mirroring and Analyzers in an EVPN-VXLAN Environment.] • EVPN-VXLAN Layer 2 gateway, including: • Multihoming • Address Resolution Protocol (ARP) suppression • Layer 3 IPv4 underlay with integrated routing and bridging (IRB) and LAG • Core isolation • Broadcast, unknown unicast, and multicast (BUM) traffic forwarding by ingress replication only • MAC move limits [See Understanding EVPN with VXLAN Data Plane Encapsulation, EVPN Proxy Arp and Arp Suppression, and Proxy NDP and NDP Suppression, IP Fabric Underlay Network Design and Implementation, overlay-ecmp, Edge-Routed Bridging Overlay Design and Implementation, Layer 2 Interface Status Tracking and Shutdown Actions for EVPN Core Isolation Conditions, and mac-move-limit.]

(Continued)

Feature	Description
Features optimized for AI-ML fabrics	- Support for the following features: - Inband Flow Analyzer (IFA) 2.0 transit node [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.] - Test access point (TAP) aggregation [See TAP Aggregation for Network Monitoring.] - Priority-based flow control (PFC) using Differentiated Services code points (DSCPs) at Layer 3 for untagged traffic [See DSCP-based PFC for Layer 3 Untagged Traffic.] - Priority-based flow control (PFC) X-ON threshold and per-queue alpha values [See Customize PFC X-ON Threshold and Per-Queue Alpha Values.] - Data center quantized congestion notification (DCQCN) [See Data Center Quantized Congestion Notification (DCQCN).] - Priority-based flow control (PFC) watchdog [See PFC Watchdog.] - Enhanced transmission selection (ETS) [See Understanding CoS Hierarchical Port Scheduling (ETS).] - SNMP and telemetry for priority-based flow control (PFC), explicit congestion notification (ECN), and CoS ingress packet drops due to ingress port congestion [See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved, show snmp mib, and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)>. For sensors, see Junos YANG Data Model Explorer.] - Extended sFlow monitoring functionality to export sFlow sample packets through the mgmt_junos interface and non-default virtual routing and forwarding (VRF) WAN ports

(Continued)

Feature	Description
	[See collector, show sflow collector, and System Logging and Routing Instances.] Increased global shared buffer pool [See Increase Shared Buffer Pool by Reducing Dedicated Buffer.]
Firewall Filters	Support for firewall filters. [See Firewall Filter Match Conditions and Actions.]
Layer 2 and Layer 3 Multicast	- Support for the following Layer 2 and Layer 3 multicast features: - Support for Multicast Listener Discovery (MLD) snooping and integrated routing and bridging (IRB) stitching [See Understanding MLD Snooping.] - Support for multicast forwarding [See Multicast Overview.] - Support for IGMP snooping [See IGMP Snooping Overview.] - Support for PIM source-specific multicast (PIM SSM) and PIM sparse mode (PIM SM) [See PIM Overview.]
Layer 2 features	Support for Layer 2 bridging and forwarding along with control and discovery protocols such as xSTP, LACP, and LLDP. [See Spanning Tree Protocol, Aggregated Ethernet LACP for Switches, and Device Discovery Using LLDP.]

(Continued)

Feature	Description
Layer 3 features	• Support for the following L3 unicast forwarding features: • IPv4 BGP • IPv4 MBGP • BGP 4-byte autonomous system number (ASN) support • BGP Add Path (supporting 128 paths) • BGP Add Path community-based selection • IS-IS • IPv6 OSPFv3 • IPv6 IS-IS • IPv6 multi-protocol BGP • IPv4 and IPv6 static routing • 128-Way ECMP • IPv4 over GRE • Graceful restart support for BGP, IS-IS, and OSPF • Path MTU discovery • IPv6 neighbor discovery protocol (NDP) • IPv6 stateless address autoconfiguration (SLAAC) • IPv6 ping • Virtual router (VRF-lite)—IS-IS, OSPF, and BGP
Load Balancing	• Support for dynamic load balancing (DLB). However, DLB is not supported on LAG. [See Dynamic Load Balancing.]

(Continued)

Feature	Description
MPLS features	• Support for the following features: • Source Packet Routing in Networking MPLS (SPRING MPLS) • Static and dynamic label-switched paths (LSPs) • LDP and RSVP • RSVP traffic engineering • Penultimate-hop popping • MPLS with ECMP • MPLS over LAG • L3 VPN for IPv4 and IPv6 • CoS within MPLS networks [See MPLS Overview, Understanding CoS MPLS EXP Classifiers and Rewrite Rules.]

(Continued)

Feature	Description
Network Management and Monitoring	- Support for sFlow and sFlow monitoring over VXLAN. [See Overview of sFlow Technology and Configure sFlow for EVPN-VXLAN Networks.] - Support for analyzers and port mirroring across Layer 2 and Layer 3 interfaces. Includes support for the following traffic monitoring and mirroring capabilities: - Native analyzer - Local port mirroring - Remote port mirroring - Remote port mirroring to an IP address using GRE encapsulation - Port mirroring support in EVPN-VXLAN environments 1:N port mirroring for sending a source packet to multiple Layer 2 destinations [See Understanding Port Mirroring and Analyzers, How to Configure Remote Port Mirroring for EVPN-VXLAN Fabrics, and 1:N Port Mirroring to Multiple Destinations on Switches.]
Protection against DDoS attacks	Support for distributed denial of service (DDoS) protection. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]

(Continued)

Feature	Description
Software installation and upgrade	• Support for zero-touch provisioning (ZTP) over IPv6 on the management and WAN interfaces. [See Zero Touch Provisioning.] • Firmware upgrade support. [See request system firmware upgrade (Junos OS Evolved).] • Secure boot and common BIOS support. 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 Overview.]

New QFX Series Liquid-Cooled Switch (QFX5250-64OE-L)—The Juniper Networks® QFX5250-64OE-L switch is a fixed-configuration, liquid-cooled device with 64 octal small form-factor pluggable (OSFP) ports that support speeds of up to 1600 gigabits per second (Gbps). Features such as 102.4 terabits per second (Tbps) throughput and 2OU shallow buffer design make this switch optimal as an end-of-row, leaf, or spine device in IP fabric architectures. The switch is compliant with and powered by Open Rack v3 (ORv3) and has a maximum power output of 4300 watt (W).

To install the QFX5250-64OE-L switch and perform initial configuration, routine maintenance, and troubleshooting, see the [QFX5250 Switch Hardware Guide](#).

See [Feature Explorer](#) for the complete list of features for any platform.

Table: QFX5250-64OE-L Feature Support

Feature	Description
Authentication and access control	• Support for 802.1X authentication with unified in-service software upgrade (ISSU). [See 802.1X Authentication.]

(Continued)

Feature	Description
Class of service (CoS)	• Support for CoS features on Layer 2 and Layer 3 interfaces, including: • IPv4 and IPv6 unicast routing • Classification and rewrite rules for Differentiated Services code point (DSCP) and IEEE-802.1p • Port scheduling • Dynamic interface-specific threshold profiles for shared buffer pools • 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. • Data Center Bridging Capability Exchange (DCBX) protocol support for PFC. • Weighted random early detection (WRED) and drop congestion notification (DCN) • Telemetry support for CoS queue statistics exported using the sensor / <code>junos/system/linecard/qmon-sw/</code> [See CoS Overview, Dynamic Threshold Profiles for Shared Buffer Pools, Understanding DCBX, priority-flow-control, and Drop Congestion Notification (DCN).] • Support for per-queue accounting of explicit congestion notification (ECN) packets. Counters on ECN-enabled queues increment when the queues experience congestion or receive packets that encountered congestion on another device. You can view per-queue ECN accounting statistics through the <code>show interfaces queue</code> command. [See CoS Explicit Congestion Notification, and show interfaces queue.]

(Continued)

Feature	Description
Features optimized for AI-ML fabrics	• Support for priority-based flow control (PFC) watchdog. [See PFC Watchdog, and congestion-notification-profile.] • BGP support for global load balancing (GLB), and GLB multilink support on IP fabric. [See global-load-balancing, and Configure GLB on 3-CLOS IP Fabric with Multilinks.] • Support for reactive path rebalancing. [See Reactive Path Rebalancing.] • Support for Inband Flow Analyzer (IFA) 2.0 transit node. [See Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring.] • Support for increased global shared buffer pool of up to 227 MB. [See Configuring Ingress and Egress Dedicated Buffers.] • Dynamic explicit congestion notification (ECN) support for automatic adjustment of threshold settings based on real-time conditions such as queue size and traffic patterns to enhance responsiveness and management during network congestion. [See CoS Explicit Congestion Notification, and Example: Configuring Static and Dynamic ECN.] • Extended sFlow monitoring support to export sFlow sample packets through the mgmt_junos interface and non-default virtual routing and forwarding (VRF) WAN ports. [See collector, show sflow collector, and System Logging and Routing Instances.] • Support for remote port mirroring to IPv4 or IPv6 addresses with DSCP, source-address, and rate-limiting parameters. 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. This functionality is sometimes referred to as extended remote port mirroring. The mirrored packets are GRE-encapsulated. [See Port Mirroring and Analyzers.]

(Continued)

Feature	Description
	SNMP and telemetry support for priority-based flow control (PFC), explicit congestion notification (ECN), and CoS ingress packet drops due to ingress port congestion. [See SNMP MIBs and Traps Supported by Junos OS and Junos OS Evolved, show snmp mib, and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface). For sensors, see Junos YANG Data Model Explorer.]
	Support for configuring per-queue alpha value to limit the buffer each queue can consume from the shared pool. [See buffer-dynamic-threshold.]
	Support for priority-based flow control (PFC) using Differentiated Services code points (DSCPs) at Layer 3 for untagged IPv6 traffic. DSCP-based PFC is required to support Remote Direct Memory Access over converged Ethernet version 2 (RoCEv2). [See Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.]
	Telemetry support for streaming IPv4 and IPv6 transit statistics using the native resource path <code>/junos/system/linecard/interface/traffic</code> or the OpenConfig resource path <code>/interfaces/interface/</code>. [See Junos YANG Data Model Explorer, and route-accounting.]
	Support to selectively enable or disable dynamic load balancing (DLB). [See rdma-opcode, dynamic-load-balance-selective, and egress-quantization.]
	Support for PFC X-ON threshold. [See xon (Input Congestion Notification).]
	Support for the <code>no-queue-pair</code> configuration to exclude queue pair from the hash calculation. [See enhanced-hash-key.]

(Continued)

Feature	Description
Forwarding and sampling	- Support for dynamic load balancing (DLB) with a configurable flowset table in DLB flowlet mode with up to 32,000 entries distributed across 128 DLB ECMP groups. [See Dynamic Load Balancing, and Configure Flowset Table Size in DLB Flowlet Mode.] - Support for BGP link-bandwidth communities to handle unequal load balancing. [See BGP Link-Bandwidth Community.] - Support to configure interfaces on your switch to use cut-through forwarding mode for faster packet forwarding and optimized packet processing. By default, the switch uses store-and-forward mode to forward packets. [See Configuring Forwarding Mode on Switches.]
Layer 2 features	- Support for Layer 2 (L2) unicast forwarding and VRRP. [See Understanding VRRP.] - Support for IGMP snooping features, including: - IGMPv1, IGMPv2, and IGMPv3 - IGMP proxy - IGMP querier at L2 - Any-source multicast (ASM) and source-specific multicast (SSM) modes - Virtual router (VRF-lite) support - Integrated routing and bridging (IRB) support [See IGMP Snooping Overview, Multicast Overview, and Integrated Routing and Bridging.]

(Continued)

Feature	Description
Layer 3 features	• Support for Layer 3 (L3) unicast forwarding and GRE tunneling. [See Generic Routing Encapsulation (GRE).] • Support for L3 multicast forwarding, including: • Physical Interface Module (PIM) first-hop router rendezvous point (RP) functionality • Multicast Source Discovery Protocol (MSDP) • Make-before-break (MBB) support for multicast receivers on existing L3 aggregated Ethernet (ae.x) or LAG interfaces. Support includes member addition, member deletion, and managing link up, and link down events. • PIM source-specific multicast (SSM) • PIM sparse mode (PIM SM) • PIM dense mode (PIM DM) • L3 multicast forwarding on integrated routing and bridging (IRB) interfaces functionality, including: • IPv4 and IPv6 multicast • IGMPv1, IGMPv2, and IGMPv3 • Multicast Listener Discovery (MLD) versions 1 and 2 • Any-source multicast (ASM) and source-specific multicast (SSM) modes [See Multicast Routing Protocols, and PIM Overview.] • Support for DHCP stateless relay on IRB interfaces and bridge domains for DHCPv4 and DHCPv6. [See DHCP Relay Agent.]

(Continued)

Feature	Description
Network management and monitoring	- Support for sFlow and egress inline sFlow with egress port (EP) recirculation and egress flex sampler. [See Overview of sFlow Technology, and sflow.] - Support for port mirroring and analyzers. The QFX5250-64OE-L switch can support a maximum of seven port mirroring sessions. [See Understanding Port Mirroring and Analyzers.] - Support for dropped-packet notification, also known as mirror on drop (MoD), to debug packet drops in real time. [See Dropped Packet Notifications to Aid in System Performance Tuning.] - Support for test access point (TAP) aggregation with ingress access control list (ACL) user-defined fields (UDF) filtering on tap ports. TAP aggregation provides many-to-many packet replication, enabling you to capture different types of data in real time for quick network monitoring. You can apply ACL with UDF on TAP interfaces to filter the traffic to be sent to the analysis tool interfaces. [See TAP Aggregation for Network Monitoring, and TAP Aggregation Enhancements.] - Support for storm control to monitor traffic levels and drop broadcast, multicast, and unknown unicast packets when traffic exceeds a specified level (storm control level). The storm control feature prevents packet proliferation and network degradation. [See Understanding Storm Control.]
Platform and infrastructure	- Support to configure firewall filters and interfaces programmatically using the Juniper Extension Toolkit (JET) APIs. [See Overview of JET APIs.] - Resiliency feature support for platform components such as power supply units (PSUs) and interfaces on the QFX5250-64OE-L device. This feature monitors component health, alerts you of errors, and takes appropriate action to restore normal operation based on error severity. [See QFX5250 Switch Hardware Guide, and Chassis-Level User Guide.]

(Continued)

Feature	Description
Protection against DDoS attacks	Support for configuration and installation of policers at the Packet Forwarding Engine level for defense from distributed denial-of-service (DDoS) attacks. By default, DDoS protection is enabled for many protocols on the QFX5250-64OE-L switch. [See Configuring Control Plane DDoS Protection Aggregate or Individual Packet Type Policers, show ddos-protection statistics, and show ddos-protection version.]
Routing policy and firewall filters	Firewall filter support on Layer 2 and Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions and Configuring Enhanced Egress Firewall Filters.]
Software installation and upgrade	- Support for zero-touch provisioning (ZTP). [See Zero Touch Provisioning.] - Support for secure zero-touch provisioning (SZTP). [See Secure Zero Touch Provisioning.] - Secure boot and common BIOS support. 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 Overview.] - Firmware upgrade support, including the new option for the leak detection controller. [See request system firmware upgrade (Junos OS Evolved).]

New QFX5241-32OD ORv3-Compliant Switch (QFX5241-32OD-DO-T1)—Starting in Junos OS Evolved Release 25.2X100-D21, we introduce the HPE Juniper Networks® QFX5241-32OD-DO-T1 that are ORv3-compliant. The QFX5241-32OD-DO-T1 switch is a fixed-configuration device with 32 octal small form-factor pluggable (OSFP) ports that support speeds of up to 800 Gigabit Ethernet (GbE). The switch is installed as a switch-tray assembly in ORv3 racks. Features such as 25.6 terabits per second (Tbps) throughput and a 1-U shallow-buffer design enable deployment as an

end-of-row, leaf, or spine device in IP fabric architectures. The switch supports 2400-watt (W) DC power supply units (PSUs) and front-to-back airflow.

To install the QFX5241-32OD-DO-T1 switch-tray assemblies in ORv3 racks 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.

New QFX5241-64OD and QFX5241-64QD ORv3-Compliant Switches (QFX5241-64OD-DO-T2 and QFX5241-64QD-DO-T2)—Starting in Junos OS Evolved Release 25.2X100-D21, we introduce the HPE Juniper Networks® QFX5241-64OD-DO-T2 and QFX5241-64QD-DO-T2 switches that are ORv3-compliant. The QFX5241-64OD-DO-T2 and QFX5241-64QD-DO-T2 switches are fixed-configuration devices with 64 octal small form-factor pluggable (OSFP) or QSFP-DD ports that support speeds of up to 800 Gigabit Ethernet (GbE). The switches are installed as switch-tray assemblies in ORv3 racks. Features such as 51.2 terabits per second (Tbps) throughput and a 2-U shallow-buffer design enable deployment as end-of-row, leaf, or spine devices in IP fabric architectures. The switches support 3000-watt (W) DC power supply units (PSUs) and front-to-back airflow. The QFX5241 platforms also support Secure BIOS and Secure Boot.

To install the QFX5241-64OD-DO-T2 and QFX5241-64QD-DO-T2 switch-tray assemblies in ORv3 racks 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.

Chassis

Support for 16T fixed platform with 24 400GbE ports and 8 800GbE ports (QFX5140-24CD8O)—

You can deploy cost-optimized, high-bandwidth fabrics on the QFX Series 1-U fixed-platform switch with a 16T switching ASIC, 24x400GbE QSFP112 ports, and 8x800GbE OSFP ports. Select flexible port configurations (1x800GbE, 1x400GbE, 2x400GbE, 4x100GbE, 2x200GbE, 4x50GbE, 4x25GbE) and 10GbE/1GbE SFP+ in-band management. Use digital optical monitoring (DOM), direct attach copper (DAC) autonegotiation, and zero-touch provisioning (ZTP) over IPv6. Precision timing through Transparent Clock (TC) and Boundary Clock (BC), secure boot with TPM 2.0, redundant power supply units (PSUs), and N+1 fans boost device resilience.

Support for fixed-chassis platform with logical FRUs and secure ZTP over IPv6 (QFX5250-64OE-L)

—Deploy the fixed-form platform as a single logical FPC or PIC system. Manage field-replaceable units (FRUs), sensors, fans, and PSUs through Linux hwmon/sysfs and platform FRU policies to meet a 6-minute boot target. Boot from a unified ISO or USB image. Identify the platform by using sysfs and /var/run/chassis/platform and use ZTP or secure zero-touch provisioning (SZTP) over IPv6 on management and revenue ports. For maintenance, issue the `request system firmware upgrade` command to update:

- BIOS
- field-programmable gate arrays (FPGA)

- solid-state drives (SSDs)
- optics or fan complex programmable logic devices (CPLDs)
- PSUs
- digital phase-locked loop (DPLL)
- Packet Forwarding Engine.

Use short or long reset presses to trigger a warm reset or BIOS swizzle.

Automatic shutdown of empty ports (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, QFX5700, and QFX5700E)—The QFX5K Series switches support automatic shutdown of empty ports to reduce power consumption by disabling SerDes on ports without transceivers. The system monitors transceiver insertion and removal events. The system seamlessly reactivates ports when you insert transceivers and disables them again when you remove transceivers or mark them unused. You do not need additional configuration. The process does not affect active ports or disrupt network operation.

Note: On QFX5700 and QFX5700E platforms, this feature is not supported when the switch is operating in FIPS mode.

[See [Automatic Shutdown of Empty Ports](#)]

Class of Service (CoS)

Support for latency-based ECN (QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-64OD, and QFX5241-64QD)—You can use latency-based explicit congestion notification (ECN) to mark congestion based on instantaneous per-packet latency instead of buffer occupancy, reducing loss and optimizing utilization. Configure a mandatory per-queue threshold in nanoseconds, from 1-1000000000 (1ns-1s), using [set class-of-service schedulers *scheduler-name* latency-ecn threshold *threshold-value*]. Verify status with [show class-of-service scheduler-map]. Use per-port ECN counters to monitor marked traffic. Use latency-based ECN for unicast traffic only. Latency-based ECN cannot coexist with buffer-based ECN on the same port or queue.

[See [Latency-Based ECN](#).]

UEC-compliant DSCP-based packet trimming for L3 packets with trimmable-to-trimmed mapping (QFX5240-64OD, QFX5240-64QD, QFX5241-64OD, QFX5241-64QD, and QFX5241E-64OD)—Use DiffServ code point (DSCP)-based marking to perform Ultra Ethernet Consortium (UEC)-compliant packet trimming on Layer 3 traffic, including Ultra Ethernet Transport (UET), without a drop congestion notification (DCN) shim header to meet standards and manage congestion. Configure trimmable and trimmed DSCP categories and per-device mappings (such as one->four, two->five, three->six, and so on) so only trimmable packets are trimmed and remarked with the associated

trimmed code point. A DSCP rewrite applied to the DCN forwarding class takes precedence over other forwarding classes. Choose either DSCP-based trimming or DCN-header-based trimming—the two modes are mutually exclusive.

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

EVPN

IGMP snooping in EVPN-VXLAN BO networks with IPv4 or IPv6 underlays (QFX5130-32CD, QFX5130-48, QFX5130-48CM, and QFX5700)—With IGMP snooping, the devices in an EVPN-VXLAN network forward multicast traffic flows only to the hosts that subscribe to those flows, rather than flooding the traffic to all hosts. Now you can enable L2 multicast with IGMP snooping for intra-VLAN IPv4 multicast traffic in an EVPN-VXLAN Bridged Overlay (BO) network with:

- MAC-VRF EVPN instances with `vlan-based` or `vlan-aware` service types.
- Enterprise-style interfaces.

You must configure the bridge domains (VLANs) symmetrically on EVPN multihoming peer devices.

[See [Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment](#).]

Copy and preserve 802.1p CoS priority bits between VLAN S-tags and C-tags with EVPN-VXLAN (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, and QFX5700E)—

Use this feature to copy and preserve the 802.1p priority between customer VLAN headers (C-tags) and service provider VLAN headers (S-tags) in EVPN-VXLAN traffic. We support this feature with service provider-style interface configurations. If you enable this feature:

- When a leaf device adds an S-tag to a packet, this feature copies the CoS priority bits from the C-tag to the S-tag.
- When a leaf device removes an S-tag from a packet, this feature retains the priority bits from the C-tag.
- The traffic priority remains unchanged in the VXLAN tunnel.

Use the `vxlan-enable-dot1p-copy` option at the `[edit forwarding-options]` hierarchy level to enable this feature. This setting overrides VLAN rewrite settings.

[See [Copy 802.1p Class of Service Priority from VLAN C-tags to VLAN S-tags in EVPN-VXLAN Network Traffic](#) and `vxlan-enable-dot1p-copy`.]

Multihoming peers in EVPN-VXLAN bridged overlays automatically synchronize multicast router interface status (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, and QFX5700)—Multihoming

peer EVPN-VXLAN border leaf (BL) devices that connect to external multicast receivers will now automatically synchronize the multicast router (mrouter) interface status for the EVPN segment identifiers (ESIs) the peer devices share. This feature ensures that multicast traffic reaches external receivers when a non-designated forwarder (NDF) BL peer device receives IGMP or Multicast Listener Discovery (MLD) queries from the external PIM domain.

The peer BL devices:

- Use EVPN Type 7 Join Sync routes to communicate the mrouter interface status among their ESI peers.
- Advertise EVPN Type 3 Inclusive Multicast Ethernet Tag (IMET) routes that enable the other EVPN devices to learn how to reach the mrouter interfaces.

This behavior happens by default on supported devices, removing the need to manually configure the `multicast-router-interface` setting on the multihoming peer BL devices.

We support this feature in EVPN-VXLAN bridged overlay (BO) networks with IPv4 or IPv6 underlay peering.

[See [Multicast Router Port Status Synchronization in EVPN-VXLAN Bridged Overlays](#).]

Symmetric IRB for IPv6 underlay with EVPN Type-2 MAC-IP routes for intersubnet routing (QFX5130-32CD and QFX5130-48C)—You can now use symmetrical integrated routing and bridging (IRB) to forward traffic consistently across all provider edge (PE) devices. The device carries EVPN Type 2 MAC-IP routes across an IPv6 underlay that provides ECMP-based transport for scalable, multitenant Layer 2 (L2) and Layer 3 (L3) connectivity. For each routing instance, use the `irb-symmetric-routing vni` statement under the `[edit protocols evpn]` hierarchy level to maintain symmetrical bridging and routing.

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

Egress Rate Limiting for VXLAN Tunnel-Initiated Traffic per VNI (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, and QFX5700E)—You can enforce per-VNI egress rate limits on VXLAN tunnel-initiated traffic to prevent congestion, mitigate denial-of-service (DoS) risk, and prioritize critical services. This configuration targets traffic entering the VXLAN tunnel. This does not rate limit locally switched or routed traffic. We added a new egress VLAN access control list (ACL) filter profile to support this feature. You enable this feature by configuring the `egress profile1` filter profile, a filter with the `vxlan tunnel-initiated` match qualifier and a policer with the `discard` action. You attach the filters per VLAN by configuring `set routing-instances <instance-name> vlans <vlan-name> forwarding-options filter output <filter-name>`. You use `show firewall` to view the policer statistics.

[See [ethernet-switching \(firewall profiles\)](#) and [tunnel-initiated](#).]

Forwarding Options

Symmetric hashing support (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, QFX5241E-64OD, and QFX5700)—Symmetric hashing ensures that bidirectional traffic flows use the same path across Equal-Cost Multi-Path (ECMP) routes or Link Aggregation Groups (LAGs). This behavior is important for services that require flow symmetry, such as NAT, firewalls, DDoS protection, and data collection.

When enabled, symmetric hashing normalizes Layer 3 and Layer 4 fields before computing the hash. This normalization guarantees that forward and reverse traffic produce the same hash value, mapping to the same link or ECMP path.

To enable symmetric hashing, use the following configuration statement:

- `set forwarding-options enhanced-hash-key symmetric-hash inet/inet6`

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

High Availability (HA) and Resiliency

Inline micro BFD support (QFX5130-32CD, QFX5130E-32CD, QFX5700, and QFX5700E)—Inline micro BFD sessions can be used to monitor the status of individual member links in a LAG bundle. Use the `set interfaces interface aggregated-ether-options bfd-liveness-detection` configuration statement to configure micro BFD.

The BFD local address is the loopback address of the source of the micro BFD sessions. Micro BFD sessions do not ISSU or authentication.

[See [Understanding Independent Micro BFD Sessions for LAG](#).]

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 [Configuring BFD](#).]

Interfaces

Support for mixed-speed and breakout on 24 x 400G QSFP112 and 8 x 800G OSFP800 ports (QFX5140-24CD80)—Use QFX5140-24CD80 switches to deploy 400G/800G switching with 24 x 400G QSFP112 and eight x 800G OSFP800 front-panel ports. Configure mixed port speeds from 25G to 800G. Apply breakouts (2x50G, 4x50G, 4x100G, and 2x200G) to tailor bandwidth. Plan

long-reach (LR) interconnects by selecting appropriate optics to extend connectivity across data center sites. See [Port Speed on QFX Series Switches](#).

Support for native 1600G ORHS with 2x/4x/8x channelization and 25G/10G SFP28

(QFX5250-64OE-L)—Configure QFX5250-64OE/OE-L ports for native 1600 Gbps/800 Gbps on OSFP-RHS (ORHS) and 25 Gbps/10 Gbps on SFP28. Channelize ORHS ports into 2x [800/400/200], 4x [400/200/100], or 8x [200/100/50]. Use this flexibility to align interface speeds with workload requirements. See [Port Speed on QFX Series Switches](#).

Support for Enhanced PRBS testing (QFX5220, QFX5240-64OD, and QFX5240-64QD)—Use the PRBS test on the QFX5K device ports to validate signal integrity and assess link quality before deployment. Use the test chassis prbs fpc <fpc num> pic-slot <pic num> port <port num> direction <dir type> pattern <num> command to start a test. You can select the speed or the polynomial (P7-P31) options to run the test. The polynomial option provides the size of the bit sequence that is sent over a network link or a transceiver. Set the direction option to determine whether the port acts as a PRBS generator or a PRBS checker. Execute the clear interface statistics command on the port before testing. Use the show interfaces extensive command to display the PRBS state and the link error statistics.

[See [show interfaces extensive](#) and [gNOI Diagnostic \(Diag\) Service](#).]

Network Management and Monitoring

Support for full-length SNMP engine ID (QFX5220)—Use the unique-id <hex-string> command in the engine-id configuration statement to enable a full-length SNMP engine ID. The unique ID uses a hexadecimal string of the length between 5 and 32 bytes.

[See [engine-id](#), [Configure SNMP in Junos OS](#), and [Configure Engine ID on SNMPv3](#).]

Ingress inline sFlow and sFlow collector statistics (QFX5130-32CD, QFX5130-48CM, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E.)—Ingress inline sFlow samples packets on ingress local ports in real time, providing timely visibility into live traffic without relying on control plane processing. The device handles sampling decisions, statistics collection, and sample-destination selection in the Packet Forwarding Engine after ingress processing and before queuing, minimizing forwarding impact. The device copies selected packet and mirrors to a configured analyzer (mirror-to-port) for detailed analysis while traffic moves normally. Inline operation ensures accurate flow insight at high speeds, with overhead controlled by a configurable sampling rate and consistent behavior regardless of the egress port.

[See [Ingress Inline sFlow Overview](#).]

Support for PBLQ testing using gNOI LinkQualification service RPCs (QFX5220, QFX5240-64OD, and QFX5240-64QD)—You can execute supported LinkQualification service remote procedure calls (RPCs) to perform a packet-based link qualification (PBLQ) test on a pair of connected ports. A PBLQ test returns information about the quality of the physical link with statistics such as the number of packets sent, received, and dropped. The client device runs the test by injecting

packets into one device and remotely reflecting them onto the peer. You can then manage and load-balance traffic across the network based on the reported link quality. Using gRPC Network Operations Interface (gNOI) operations like this one enables you to use the same suite of microservices to efficiently manage large-scale multivendor networks.

To monitor the status of the PBLQ test, use the `show services monitoring pblq (capabilities | summary | test-id test-id)` command.

[See [gNOI LinkQualification Service](#), [processes](#), and [show services monitoring pblq](#).]

sFlow sampling with firewall filters (QFX5220, QFX5230-64CD, QFX5240-64OD, and QFX5240-64QD)—Junos OS Evolved supports sFlow packet sampling using ingress firewall filters. This feature enables selective packet sampling of traffic that matches specific firewall filter terms by using the `sample` action, providing granular traffic visibility while minimizing overall sampling overhead.

You configure packet sampling in a firewall filter term by using the `set firewall family inet filter filter-name term term-name then sample` command. You must combine this configuration with standard sFlow configuration under the `set protocols sflow` hierarchy, such as `sample-rate` and `collector` for configuring sFlow.

When you configure the `sample` action in a firewall filter term, you cannot configure the `discard` or `reject` action in the same term.

[See [Filter-Based sFlow Packet Sampling Overview](#).]

DC Flow Analyzer (QFX5240-64OD and QFX5240-64QD)—We support a new DC Flow Analyzer (DFA) feature, which provides dynamic, per-flow visibility on the switch. The DFA also helps track RDMA ROCEv2 flow destination queue pair information in Junos OS Evolved and other AI DC workloads. The DFA learns flows over time based on a default tuple (source IP, destination IP, IP protocol version, ingress interface, Layer 4 source port, Layer 4 destination port, and destination queue pair), and tracks counters for each learned flow. The DFA monitors both forwarded and dropped traffic and supports IPv4 and IPv6 flows. You can selectively learn flows by applying filters (for example, matching on destination IP address or UDP destination port) and group flows for targeted monitoring. To enable this feature, configure the `flow analyzer` statement at the `[edit]` hierarchy level and specify the DFA profile name as an action for a firewall filter. To view information about the learned flows, use the `show flow table` command.

[See [DC Flow Analyzer](#).]

Stateless flow latency monitoring (QFX5240-64OD and QFX5240-64QD)—Monitors packet residence time—travel time from input interface to output interface—on the switch and mirrors packets whose latency falls within the residence-time range that you define. Matching packets are exported in IPFIX format with latency and flow metadata to an external collector for analysis,

without maintaining per-flow state or impacting forwarding. Stateless flow latency monitoring mirrors all packets in the configured range to the collector.

Precision Time Protocol (PTP)

Transparent clock support (QFX5140-24CD8O)—QFX5140-24CD8O switches support the Precision Time Protocol (PTP) transparent clock feature. Transparent clocks improve synchronization between the timeTransmitter and timeReceiver clocks and ensure that the timeTransmitter and timeReceiver clocks are not impacted by the effects of packet delay variation.

[See [PTP Transparent Clocks](#).]

Routing Protocols

Path quality profile sharing for GLB multilink support on IP fabric (QFX5240-64OD and QFX5240-64QD)—While we support larger Clos networks and more graphic processing cards (GPUs), these QFX Series switches support only 64 profiles. In Clos networks with five or more stages, some nodes, such as super spines, exceed 64 next-next-hop nodes. We can reuse the profiles under specific conditions to support more than 64 next-next-hop nodes. We support global load balancing (GLB) for Clos networks with profile sharing in hyperscaler artificial intelligence--machine learning (AI-ML) fabrics containing tens of thousands of leaf devices or GPUs.

[See [profile-sharing](#).]

Support for interface-scoped TCP authentication for BGP unnumbered peers (QFX5700)—We are extending support for interface-scoped authentication for BGP unnumbered peers in single-hop connectivity. The interface-scoped configurations enable TCP-MD5, TCP-Keychain, TCP AO, and Ron Bonica-draft authentication for BGP unnumbered peers using IPv6 link-local addresses. The authentication method is best suited for deployments with symmetric routing, where traffic for a TCP flow traverses the same logical interface in both directions. Authentication keys are associated directly with logical interfaces instead of IP addresses or prefixes as BGP unnumbered auto-discovered peering uses IPv6 link-local addresses. Authentication is applied based on the logical interface over which the BGP session is established and is independent of the dynamically learned peer address.

We also support nonstop active routing (NSR) support for interface-scoped authentication for BGP auto-discovered neighbors.

Note: When an authentication algorithm is specified, it must be configured with a key-chain and cannot be used with key.

[See <https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/interface-scoped-authentication-bgp-unnumbered-peers.html>.]

Services Applications

IFA 2.0 initiator and terminator support with new live and probe modes (QFX5240-64OD and QFX5240-64QD)—You can configure Inband Flow Analyzer 2.0 (IFA 2.0) initiator nodes to operate in either live or probe mode.

In live mode, IFA 2.0 uses the original packets as IFA probe packets. In probe mode, IFA 2.0 uses copies of the original packets as the IFA probe packets.

In live mode, the initiator inserts IFA headers and metadata into the live packets. In probe mode, the initiator samples the flow and copies that sample to create the probe.

The IFA terminator node parses the IFA header in the packet to determine if it is a live or probe packet. If it is a live packet, then the terminator strips the IFA layer and forwards the modified live packet to its destination. If the packet is a probe packet, then the terminator drops the packet, and sends a mirror copy to the collector.

To specify the mode, configure the `mode (live | probe)` statement at the `[edit services inband-flow-telemetry]` hierarchy level. The default mode is probe. Use the operational mode command `show services inband-flow-telemetry global` to see results that include the mode information.

[See [Inband Flow Analyzer \(IFA\) 2.0 Real-Time Flow Monitoring](#) and [show services inband-flow-telemetry](#).]

Additional Features

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

Hash-based PIM RPF selection (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5700, and QFX5700E)—Use hash-based Protocol Independent Multicast (PIM) reverse path forwarding (RPF) selection to route multicast traffic for a specific source and group (S,G) consistently through the same upstream node. The device uses multilevel hashing of PIM neighbor attributes such as router ID, cluster ID, and interface ID for consistent routing. To configure hash-based PIM RPF selection, include the `hash-based-rpf-selection` statement at the `[edit protocols pim]` hierarchy level.

[See [Hash-based PIM RPF Selection](#) and [hash-based-rpf-selection](#).]

Platform resiliency, ASIC fault simulation, and fru-get-state diagnostics (QFX5140)—You can deploy a 1-U top-of-rack switch with 24x400GbE, 8x800GbE, and 2x25GbE ports, supporting 25–800-Gbps port speed configuration, in leaf, spine, or super-spine roles. Leverage platform resiliency across control, switching, storage, optics, clocking, cooling, and power supply along with fault detection, alarms, and recovery procedures. Use `fru-get-state` to collect detailed device, clock, power, sensor, and PCI Express diagnostics for rapid troubleshooting. Validate fault handling with commands such as `<codeph>test brcm_xgs simulate-fault asic init-fault fpc 0`

enable/disable </code> and <code>test brcm_xgs simulate-fault asic access-fault fpc 0 enable/disable</code>.

Platform software support for second-source SSD (QFX5241-64OD)—We provide support for inbuilt routing engine, control board, power supply, fan trays, PICs, and flexible PIC concentrators (FPCs) on QFX5241-64OD switches with AC power supply units. QFX5241-64OD switches support all features relevant to fixed form factor switches. The fans are field replaceable units (FRUs), providing N+1 fan redundancy. You can manage the fan speed according to the environment monitoring (EM) policy. Two power supply units (PSUs) in a switch provides 1+1 power supply redundancy. You can manage a QFX5241-64OD switch and its components such as sensors, interfaces, and so on by using existing CLI commands. [See [temperature sensor](#) and [Configure Port Speed](#).]

Overlapping VLAN support for ERB EP-style logical interfaces (QFX5130-32CD, QFX5130-48C, and QFX5700)—These switches support overlapping customer VLAN IDs on EP-style Ethernet-Switching interfaces by scoping VLANs per MAC-VRF using forwarding-instance identifiers in EVPN-VXLAN ERB deployments.

[See [Overlapping VLAN Support Using Multiple Forwarding Instances or VLAN Normalization and forwarding-instance](#).]

Resilient hashing for EVPN-VXLAN overlay ECMP (EVPN Type 5) (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5130E-32CD, QFX5700, and QFX5700E)—This feature adds resilient hashing for overlay ECMP next hops learned from EVPN Type 5 (IP prefix) routes in EVPN-VXLAN fabrics. When an overlay ECMP member (remote VTEP/gateway) is removed, per-flow forwarding remains stable, minimizing flow remapping compared to modulo-based static hashing. When both dynamic load balancing (DLB) and resilient hashing are enabled, DLB takes precedence for overlay ECMP. This behavior extends existing resilient hashing support (non-VXLAN ECMP and VXLAN underlay ECMP) to overlay ECMP built from EVPN Type 5 routes.

[See [Resilient Hashing on LAGs and ECMP groups](#).]

Support for SRv6 network programming (QFX5240-64OD, QFX5240-64QD, QFX5241-32OD, QFX5241-32QD, QFX5241-64OD, QFX5241-64QD, and QFX5241E-64OD)—We support the following SRv6 features:

- Layer 3 Services over SRv6 in BGP (End.DT4 and End.DT6)
- Operations, Administration, and Maintenance (OAM) ping for SRv6
- SRv6 traceroute
- SRv6 header compression (uSID or micro-SID)

[See [Understanding SRv6 Network Programming in IS-IS Networks](#).]

1:N port mirroring for sending a source packet to multiple Layer 2 destinations (QFX5130-32CD, QFX5130E-32CD, QFX5130-48C, QFX5130-48CM, QFX5220, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)—Mirror traffic to multiple Layer 2 destinations. Use one or both of the following configurations:

- A port-mirroring instance that is based on a firewall filter. Use the configuration statements in the `[edit forwarding-options port-mirroring instance]` hierarchy
- A native analyzer. Use the configuration statements in the `[edit forwarding-options analyzer]` hierarchy.

For both configuration methods, configure next-hop groups to direct the mirrored packets to their destinations.

[See [1:N Port Mirroring to Multiple Destinations on Switches](#).]

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)—Apply ingress access control list (ACL) with user-defined fields (UDF) on tap interfaces. This helps you choose traffic to be sent to tool interfaces on a tap-aggregation switch. If the ACL match conflicts with the tap-aggregation rule, the ACL match takes precedence. Configure the TAP aggregation interfaces under `[edit interfaces]`:

Add interface to tap group:

```
set interfaces interface-name unit 0 mode tap group tap-group-name
```

Add interface to tool group:

```
set interfaces interface-name unit 0 mode tool group tool-group-name
```

Timestamp option for tap-aggregation packets (QFX5220-128C and QFX5220-32CD)—Use the tap-aggregation feature to insert a timestamp in packets at data capture—before the packets are sent to the tool ports for analysis. The timestamp shows exactly when the data packet was captured. Configure the Precision Time Protocol (PTP) reference clock on the tap-aggregation switch and ensure that PTP is 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 To enable the timestamp option:

```
set interfaces interface-name timestamp ingress
```

BGP support for deterministic path forwarding (DPF) in a CLOS network (QFX5130-32CD, QFX5130-48C, QFX5130-48CM, QFX5230-64CD, QFX5240-64OD, QFX5240-64QD, QFX5700, and QFX5700E)—BGP Deterministic Path Forwarding (DPF) divides a physical fabric into multiple logical fabrics so that different flows can be mapped to different logical fabrics to serve the requirements of the flows. When DPF is used for EVPN-VXLAN AI DC use cases, overlay forwarding

can be controlled using the colored path for the given RT5 EVPN instance or for the MAC-VRF (EVI).

The VXLAN tunnel can be enabled with a specific source IP address, which can be advertised with a different BGP color community and can follow a specific BGP colored peering from leaf to spine.

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 Egress Traffic Engineering](#).]

Optics Support for QFX5250 Switch (QFX5250)—Supported transceivers, optical interfaces, and DAC cables—Select your product in the Hardware Compatibility Tool ([Hardware Compatibility Tool](#)) to view supported transceivers, optical interfaces, and direct attach copper (DAC) cables for your platform or interface module. We update the HCT and provide the first supported release information when the optic becomes available.

What's Changed in 25.2X100-D20

IN THIS SECTION

- [Services Applications](#) | 30

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

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

Services Applications

- [IFA metadata stack enhancement (QFX Series)]—Additional class-of-service fields are now encoded in the Inband Flow Analyzer (IFA) metadata. Along with the existing fields, three egress-queue related counters are included: Egress Queue Tx Bytes, Egress Queue Current Bytes (Current Buffer Occupancy), and Egress Queue Maximum Bytes (Peak Buffer Occupancy).

[See [Inband Flow Analyzer \(IFA\) 2.0 Real-Time Flow Monitoring](#).]

Known Limitations in 25.2X100-D20

IN THIS SECTION

- [General Routing](#) | 30

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

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

General Routing

- On QFX5130 platforms, in VXLAN topologies, symmetric ECMP hashing might not be enforced when resilient hashing is concurrently enabled.

[PR1907397](#)

- On the QFX5230/QFX5240 platforms, host injected control protocol frames are accounted under the egress port's best-effort queue (Q0 by default) even though they are sent on the network control queue. This is only a display issue of the `show` command output and does not affect the actual prioritization of traffic.

[PR1835046](#)

- On QFX5250 platform, Global Load Balancing (GLB) may not operate as expected in certain topologies. This can impact traffic distribution and remote fail over behaviour.

PR1927407

- When a MPLS packet ingress on sub-interface or IRB interface with a vlan-tag gets dropped in TD5 ASIC ingress pipeline because of vlan membership check failure.

PR1935329

- When label action is PHP and ingress packet has only one label matching PHP action label and that packet's IP header has Router Alert (IP Options) then TD5 ASIC is not updating ether Type field in Ethernet Header of egress packet. Which may result in peer device dropping the packet.

PR1951690

Open Issues in 25.2X100-D20

IN THIS SECTION

- [General Routing | 31](#)
- [Interfaces and Chassis | 32](#)

Learn about open issues in this release for QFX Series switches.

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

General Routing

- ARP resolution failure when there is quick flap of core file facing interface on scaled setup
[PR1816344](#)
- With a back-to-back connection using 800G ZR* optics, if one end is configured with native speed (default) and other end is configured with 8x100G mode, then the expectation is that link is down on both ends because of the configuration mismatch between the ends; but one of the two ends shows as Up. The same behaviour is presented by 400G ZR* optics also.
[PR1884559](#)
- On a QFX5130 or QFX5140 platform, the traffic destined to an overlay ECMP in a Type-5 routed EVPN VXLAN topology does not load balance efficiently when the following conditions are met

1. dynamic load balancing is enabled on the underlay and (2) the interface state is flapped
[PR1953665](#)

- The request `system zeroize` command is used to erase any configuration or custom data, and reset the system to factory defaults. After running this command on QFX5250-64OE-L, the management interface MAC address is shown with the default value of `00:a0:c9:00:00:00` as displayed by the `show interfaces re0:mgmt-0 | grep Hardware address` command. This does not have any functional impact.

[PR1957903](#)

- When the congestion is at VXLAN encap node the inner packet IP ECN is getting remarked correctly. Since default behaviour is to copy the OUTER IP ECN to inner at the decap node the packet ECN is missing at the Access side of the Decap node. This can overcome by having the below CLI at decap node: `set forwarding-options vxlan-disable-copy-tos-decap`

[PR1932587](#)

Interfaces and Chassis

- Link with the DAC (SFP56-50G-DAC-3M) comes up with 25G default speed configuration when switch is rebooted. Hence, when 50G speed configuration is applied, peer side sees the link flap. [PR1836697](#)
- Encountering input-errors and framing-errors on channelizing QSFP56-DD-400GBASE-DR4.

[PR1848109](#)

- This is a limitation of the current module firmware on the 1.6T optics used with QFX5250-64OE-L. PRBS test fails when 2km Appels are used on the 1.6T optics modules.

[PR1957938](#)

- Packet drops due to continuous CRC errors are observed after changing the port speed (4x25G >> 2x200G >> 4x25G) - It is negative test case as we don't support 2X200G for optics: QSFP-100GBASE-SR4 Goal was to set some unsupported speed and revert it back (NRZ->PAM->NRZ). It should behave normal. work-around : As CDR is locked with incorrect incoming signal when the PLL changes, After Restarting the PMD, we see CRC not incrementing.

[PR1925924](#)

Migration, Upgrade, and Downgrade Instructions

For products impacted, see [Feature Explorer](#).

Follow these steps to upgrade your Junos OS Evolved software:

1. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage: [Downloads](#).
2. In the Find a Product box, enter the Junos OS platform for the software that you want to download.
3. Select Junos OS Evolved from the OS drop-down list.
4. Select the relevant release number from the Version drop-down list.
5. In the **Install Package** section, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the device or to your internal software distribution site.
10. Install the new package on the device.

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

For more information about software installation and upgrade, see [Software Installation and Upgrade Overview \(Junos OS Evolved\)](#). For more information about EOL releases and to review a list of EOL releases, see [Junos OS Evolved Dates & Milestones - Juniper Networks](#).

Documentation Updates

There are no documentation updates for this release.

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.

[Feature Explorer](#)

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

[index?page=prsearch](#).

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

[Hardware Compatibility Tool](#)

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.

Requesting Technical Support

IN THIS SECTION

- Self-Help Online Tools and Resources | 35
- Creating a Service Request with JTAC | 36

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit [Product Warranty Policy - Support - Juniper Networks](#).
- 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: [Support](#)
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: [Documentation](#)
- Find solutions and answer questions using our Knowledge Base: [CEC Juniper Community](#)
- Download the latest versions of software and review release notes: [Support](#)

- Search technical bulletins for relevant hardware and software notifications: [CEC Juniper Community](#)
- Join and participate in the Juniper Networks Community Forum: [Elevate - Elevate Community](#)

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

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: [Support](#)
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: [Documentation](#)
- Find solutions and answer questions using our Knowledge Base: [CEC Juniper Community](#)
- Download the latest versions of software and review release notes: [Support](#)
- Search technical bulletins for relevant hardware and software notifications: [CEC Juniper Community](#)
- Join and participate in the Juniper Networks Community Forum: [Elevate - Elevate Community](#)

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

Creating a Service Request with JTAC

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

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

For international or direct-dial options in countries without toll-free numbers, see [Contact - Support - Juniper Networks](#).

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

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

For international or direct-dial options in countries without toll-free numbers, see [Contact - Support - Juniper Networks](#).

Revision History

24-Jun-26 - Revision 4, Junos OS Evolved Release 25.2X100-D20.

22-May-26 - Revision 3, Junos OS Evolved Release 25.2X100-D20.

18-May-26 - Revision 2, Junos OS Evolved Release 25.2X100-D20.

15-May-26 - Revision 1, Junos OS Evolved Release 25.2X100-D20.