HPE JUNIPER networking

Data sheet

QFX5200 SWITCH

HPE

Product overview

The QFX5200 line of access switches is ideally suited for leaf deployments in next-generation IP fabrics. It features line-rate, high-density SFP28/QSFP+/ QSFP28 ports that support 10GbE, 25GbE, 40GbE, 50GbE, and 100GbE connections. Offering an advanced L2, L3, and MPLS feature set, network operators can use the QFX5200 to build large, high-density IP fabrics that support upgrades to 25GbE servers, network virtualization, and intelligent traffic forwarding based on proven, Internet-scale technology. The QFX5200 also supports a version of Junos OS software in which control plane and data plane processes and functions run in parallel, maximizing utilization of the high performance quad-core CPU.

Product description

Network operators are increasingly deploying scale-out, spine-and-leaf IP fabric architectures, built with fixed-configuration switches to support growing east-west traffic in the data center. The Juniper Networks QFX5200 line of next-generation, fixed configuration leaf/access switches offers flexible, cost-effective, high-density 10GbE, 25GbE, 40GbE, 50GbE, and 100GbE interfaces for server and intra-fabric connectivity, providing deployment versatility, investment protection, and future-proofing for today's data centers.

In addition to leaf deployments in IP fabric architectures, the QFX5200 switches are also a great fit for multi-tier, multichassis link aggregation (MC-LAG) network implementations commonly used in enterprise networks.

Like other QFX Series switches, the QFX5200 supports advanced Layer 2, Layer 3, and MPLS features. For large public cloud providers—among the first to adopt 25GbE servers to meet explosive workload growth—the QFX5200 enables very large, dense, and fast IP fabrics based on proven, Internet-scale technology. For enterprise customers seeking investment protection as they transition their server farms from 10GbE to 25GbE, the QFX5200 line provides native 10GbE and 25GbE downlink ports as well as 40GbE and 100GbE uplink connections.

The QFX5200 runs the same reliable, high performance Juniper Networks Junos operating system used by network operators around the world. The QFX5200 supports a modular version of Junos OS that allows the switch's control plane and data plane processes and functions to run in parallel, maximizing utilization of the high performance quad-core CPU.

QFX5200 switch models

The QFX5200 switches are compact, 1 U platforms that provide wire-speed packet performance, very low latency of 420ns, and a rich set of Junos OS features. In addition to a high throughput Packet Forwarding Engine (PFE), the performance of the QFX5200 control plane is further enhanced with a powerful 1.8 GHz quad-core Intel® CPU with 16 GB of memory and 64 GB SSD storage.

Two QFX5200 switch models are available:

— **QFX5200-48Y[EOL]:** A 10GbE/25GbE data center access switch that offers 48 small form-factor pluggable plus (SFP+) transceiver ports or SFP28 ports and six QSFP28 ports. The six QSFP28 ports can be configured as 6x40GbE or 6x100GbE, with an aggregate throughput of 3.6 Tbps or 4.2 Bpps per switch.

— **QFX5200-32C:** A 100GbE data center access and aggregation switch that offers 32 quad small form-factor pluggable plus (QSFP+) transceiver or QSFP28 ports, with an aggregate throughput of 6.4 Tbps or 3.2 Bpps per switch. A QSFP28 module supports a 1x100GbE connection, 2x50GbE connections, or 4x25GbE connections using breakout cables. A QSFP+ module can provide a 1x40GbE connection or 4x10GbE connections using breakout cables.

Product highlights

The QFX5200 includes the following capabilities. Please refer to the Specifications section for features currently shipping.

— **Native 25GbE configuration:** A Offering 48 ports of 25GbE and delivering 3.6 Tbps of line-rate switching, the QFX5200-48Y[EOL] is optimized for top-of-rack deployments.

— **High-density configuration:** Offering 32 ports of 100GbE or 40GbE and delivering 6.4 Tbps of line-rate switching, the QFX5200-32C is optimized for high-density fabric deployments.

— **Flexible connectivity options:** The QFX5200 offers a choice of 10GbE, 25GbE, 40GbE, 50GbE, and 100GbE interface speeds for server and intra-fabric connectivity, providing deployment versatility and investment protection.

— **Rich automation capabilities:** The QFX5200 switches support a number of network automation features for plug-and-play operations, including zero-touch provisioning, operations and event scripts, automatic rollback, and Python scripting.

— **Advanced Junos OS features:** The QFX5200 switch supports features such as BGP add-path, MPLS, L3 VPN, and Fibre Channel over Ethernet (FCoE).

— **Junos OS software architecture:** The QFX5200 supports a modular Junos OS software architecture that allows the switch's control and data plane processes and functions to run in parallel, maximizing utilization of the high performance quad-core CPU.

QFX5200 deployment options

Figure 1 shows QFX5200-48Y[EOL] 25GbE top-of-rack deployment with a Juniper Networks QFX10000 spine acting as a centralized gateway. In this topology, the QFX10000 switch can also be configured as a collapsed spine and edge device, acting as a centralized gateway for Virtual Extensible LAN (VXLAN), MPLS, and other tunneling protocols.

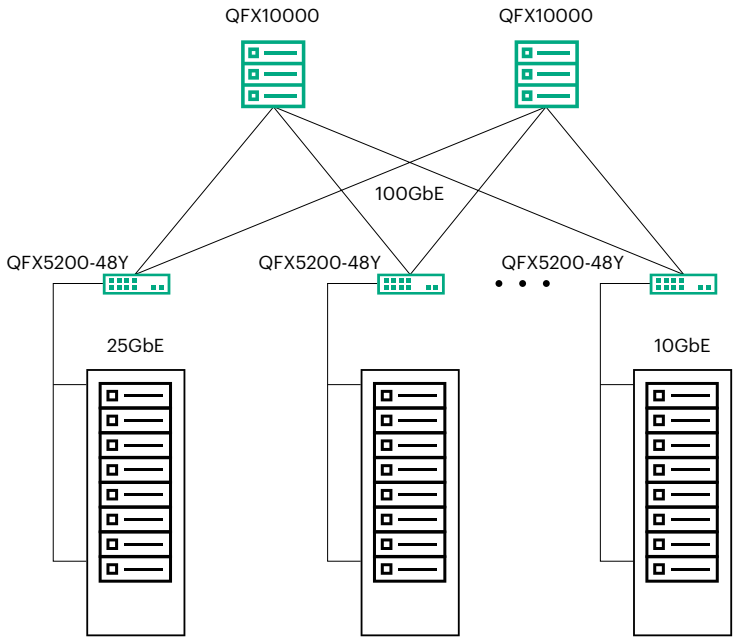

Figure 1: QFX5200-32C and QFX10000 leaf-spine deployment

The QFX5200-32C can be deployed as a top-of-rack switch in a centralized gateway architecture. Figure 2 shows the QFX5200-32C as an access switch with a Juniper Networks QFX10000 spine configured as a centralized gateway. In this topology, the QFX10000 switch can also be configured as a collapsed spine and edge device, acting as a centralized gateway for VXLAN, MPLS, and other tunneling protocols. The 100GbE ports on QFX5200-32C can also be channelized to support 4x25GbE, 2x50GbE, or 4x10GbE downlinks.

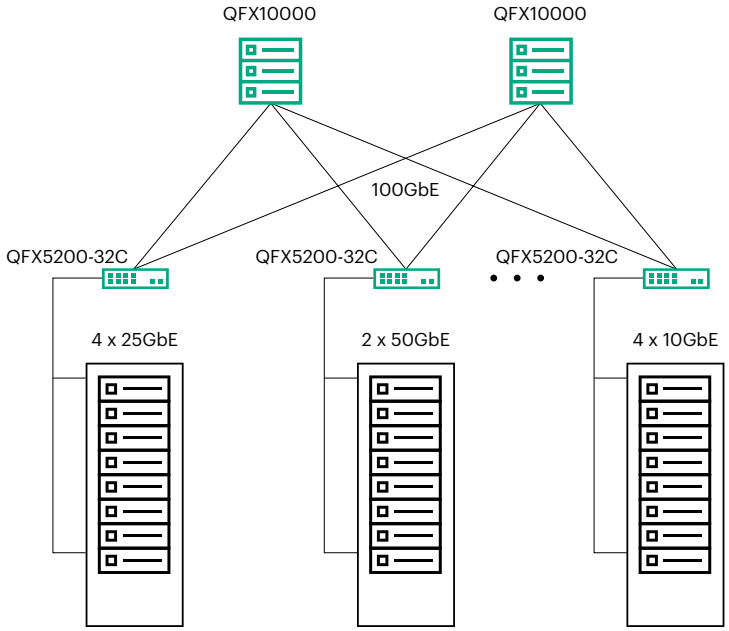

Figure 2: QFX5200-32C and QFX10000 leaf-spine deployment with centralized gateway

The QFX5200-32C can also be positioned as a spine in a distributed gateway architecture. Figure 3 shows the QFX5110 as access switches configured as a distributed gateway along with the QFX5200 spine switch. The QFX5110 switches act as a distributed gateway for VXLAN and other tunneling protocols.

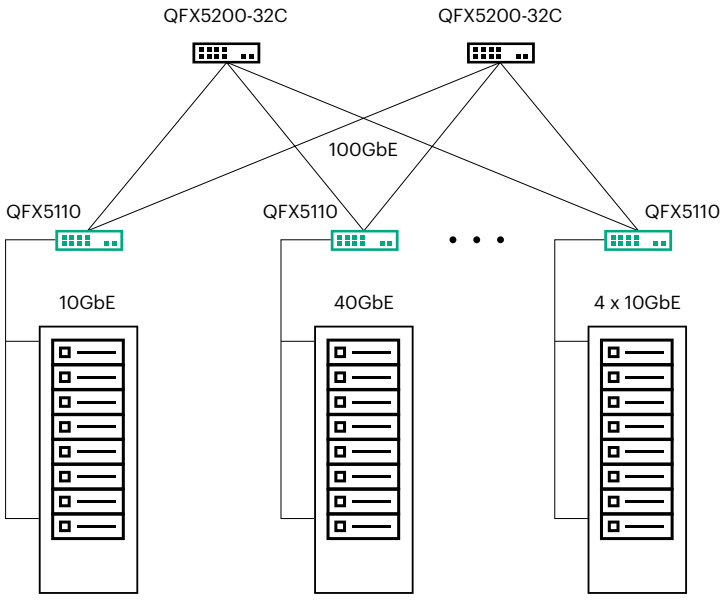

Figure 3: QFX5110 and QFX5200-32C leaf-spine deployment with distributed gateway

Architecture and key components

The QFX5200 can be used in L3 fabrics and L2 networks with MC LAG. You can choose the architecture that best suits your deployment needs and easily adapt and evolve as requirements change over time. The QFX5200 switch serves as the universal building block for these two switching architectures, enabling data center operators to build cloud networks their way.

— **Layer 3 fabric:** For customers looking to build scale-out data centers, a Layer 3 spine-and-leaf Clos fabric is ideal due to its predictable, nonblocking performance and scale characteristics. For example, a two-tier fabric built with QFX5200 switches as leafs and Juniper Networks QFX10000 switches as the spine can scale to support up to 18,432 40GbE ports, or 36,864 25GbE or 10GbE server ports in a single fabric. One of the most complicated tasks in building an IP fabric is assigning all of the implementation details, including IP addresses, BGP autonomous system (AS) numbers, routing policies, loopback address assignments, and many others. Automating the creation of an IP fabric at a large scale is equally difficult. To address these challenges, HPE has created the OpenClos project to provide free and open-source tools that automate the creation of IP fabrics in the data center. A set of Python scripts developed as an open-source project that live on GitHub, OpenClos takes a set of inputs that describe the shape and size of a data center and outputs switch configuration files, as well as a cabling plan.

— **MC-LAG:** The QFX5200 supports the MC-LAG protocol between two switches to eliminate Spanning Tree Protocol (STP) in traditional L2 networks when deployed in the aggregation layer. The active/active operation of MC-LAG ensures complete bandwidth utilization between the network's access and aggregation layers, while the dual control plane technology ensures the highest availability for applications.

Carrier-class network operating system

The QFX5200 runs the same Junos OS used by all other Juniper Networks QFX Series and EX Series Ethernet Switches, as well as the HPE routers that power the world's largest and most complex networks.

By using a common operating system, HPE delivers a consistent implementation and operation of control plane features across all products. Junos OS employs a highly available modular architecture that prevents isolated failures from bringing down an entire system. Key Junos OS features that enhance the functionality and capabilities of the QFX5200 include:

— Software modularity, with process modules running independently in their own protected memory space and with the ability to do process restarts

— Commit and rollback functionality, which ensures error-free network configurations

— A powerful set of scripts for on-box problem detection, reporting, and resolution

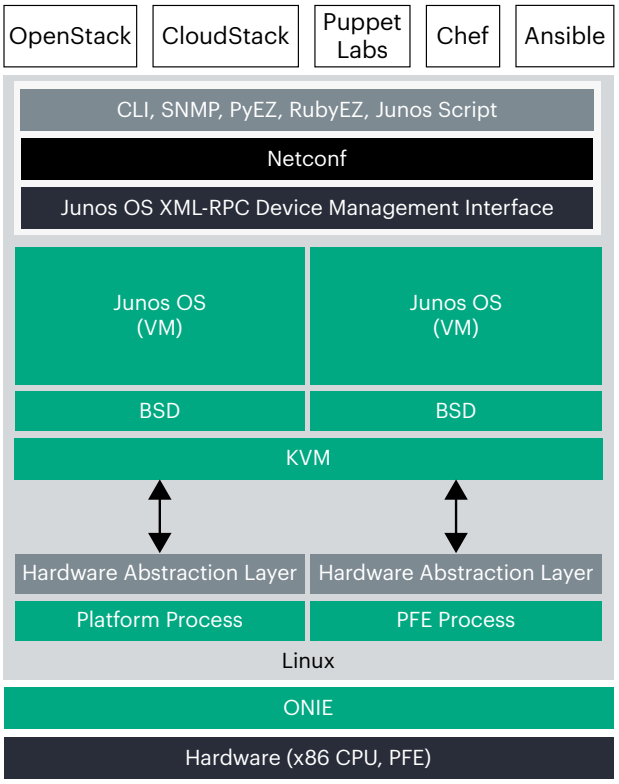

Figure 4: Junos OS software infrastructure

Junos telemetry interface

The QFX5200 supports Junos Telemetry Interface (JTI), a modern telemetry streaming tool designed for performance monitoring in complex, dynamic data centers. Streaming data to a performance management system enables network administrators to measure trends in link and node utilization, and troubleshoot such issues as network congestion in real time. JTI provides these capabilities:

— Application visibility and performance management by provisioning sensors to collect and stream data and analyze application and workload flow path through the network.

— Capacity planning and optimization by proactively detecting hotspots and monitoring latency and microbursts.

— Troubleshooting and root cause analysis via high-frequency monitoring and correlating overlay and underlay networks.

Management, monitoring, and analytics

Data center fabric management: Apstra Data Center Director (formerly Juniper Apstra) provides operators with the power of intent-based network design to help ensure changes required to enable data center services can be delivered rapidly, accurately, and consistently. Operators can further benefit from the built-in assurance and analytics capabilities to resolve day 2 operations issues quickly.

Data Center Director key features are:

— Automated deployment and zero-touch deployment

— Continuous fabric validation

— Fabric life cycle management

— Troubleshooting using advanced telemetry

For more information, see Apstra Data Center Director.

Features and benefits

— **Automation and programmability:** The QFX5200 supports numerous network automation features, including operations and event scripts, ZTP, and HPE plug-ins for OpenStack Neutron.

— **Flexible forwarding table:** The QFX5200 switch's Flexible Forwarding Table (FFT) allows the hardware table to be carved into configurable partitions of L2 media access control (MAC), L3 host, and longest prefix match (LPM) tables. In a pure L2 environment, the QFX5200 supports 136,000 MAC addresses; in L3 mode, the table supports 104,000 host entries; while in LPM mode, it supports 128,000 prefixes. A filter-based forwarding mode supports 64,000 exact match filtering rules. Junos OS provides configurable options through a CLI, enabling each QFX5200 to be optimized for different deployment scenarios.

— **MPLS:** The QFX5200 supports a broad set of MPLS features, including L3 VPN, RSVP traffic engineering, and LDP to support standards-based multitenancy and network virtualization with per-flow SLAs at scale. The QFX5200 can also be deployed as a low-latency MPLS label-switching router (LSR) or MPLS provider edge (PE) router in smaller scale environments. The QFX5200 and the Juniper Networks QFX5100 switch are the industry's most compact, low-latency, high-density, low-power family of switches to offer an MPLS feature set.

— **FCoE:** As a Fibre Channel over Ethernet (FCoE) transit switch, the QFX5200 provides an IEEE data center bridging (DCB) converged network between FCoE-enabled servers and an FCoE-to-FC gateway or FCoE-enabled Fibre Channel storage area network (SAN). The QFX5200 offers a full-featured DCB implementation that provides strong monitoring capabilities on the top-of-rack switch for SAN and LAN administration teams to maintain clear separation of management. In addition, FCoE Initialization Protocol (FIP) snooping provides perimeter protection, ensuring that the presence of an Ethernet layer does not impact existing SAN security policies. FCoE LAG support transports FCoE traffic and regular Ethernet traffic across the same link aggregation bundle over ports on a converged network adapter (CNA), while guaranteeing that the same physical link is used for an FCoE point-to-point virtual connection.

The FCoE transit switch functionality, including priority-based flow control (PFC) and Data Center Bridging Capability Exchange (DCBX), are included as part of the default software.

QFX5200-32C

QFX5200-48Y

Specifications

Hardware

Table 1: QFX5200 system capacity

	QFX5200-32C	QFX5200-48Y[EOL]
System throughput	Up to 6.4 Tbps (bi-directional)	3.6 Tbps (bi-directional)
Forwarding capacity	Up to 3.2 Bpps	4.2 Bpps
QSFP+/QSFP28 ports	32 QSFP+ or QSFP28	6 QSFP+ or QSFP28
SFP+/SFP28	128 10GbE SFP+ through QSFP+ to SFP+ breakout cable 128 25GbE SFP28 through QSFP28 to SFP28 breakout cable	48 SFP+ or SFP28
Maximum 10GbE port density	128	48
Maximum 25GbE port density	128	48
Maximum 40GbE port density	32	6
Maximum 50GbE port density	64	-
Maximum 100GbE port density	32	6

Table 2: QFX5200 system specifications

	QFX5200-32C	QFX5200-48Y[EOL]
Dimensions (W x H x D)	17.36 x 1.72 x 20.48 in. (44.09 x 4.37 x 52.02 cm)	17.36 x 1.70 x 20.28 in. (44.09 x 4.31 x 51.5 cm)
Rack units	1 U	1 U
Weight	23.8 lb (10.8 Kg)	21.2 lb (9.6 kg)
Operating system	Junos OS	Junos OS
Switch chip	Broadcom Tomahawk	Broadcom Tomahawk
CPU	Intel Quad Core Ivy Bridge 1.8 GHz CPU, 16 GB SDRAM, 64 GB SSD	Intel Quad Core Broadwell-DE 2.3 GHz CPU, 16 GB SDRAM, 32 GB SSD
Management and Precision Time Protocol (PTP) interfaces	— 1 RJ-45 1GbE port for PTP Grandmaster — 2 SMB connectors, one for pulse per second (PPS) output and one for 10 MHz clock output — 2 SFP and 1 RJ-45 10/100/1000 BASE-T for management (only two ports can be used) — 1 RS-232 console port — 1 USB 2.0 port	— 1 RJ-45 10/100/1000 BASE-T for management — 1 RS-232 console port — 1 USB 2.0 port
Power	— Redundant (1+1) hot-pluggable 850 W AC/DC power supplies — 110–240V single phase AC power — –36 to –72V DC power	— Redundant (1+1) hot-pluggable 650 W AC/DC power supplies

Table 2: QFX5200 system specifications (continued)

	QFX5200-32C	QFX5200-48Y[EOL]
Cooling	— Front-to-back and back-to-front cooling — Redundant (N+1) hot-pluggable fan modules with variable speed to minimize power draw	— Front-to-back and back-to-front cooling — Redundant (N+1) hot-pluggable fan modules
Total packet buffer	16 MB	22 MB
Warranty	HPE standard one-year warranty	HPE standard one-year warranty

Software performance scale (Unidimensional)

— MAC addresses per system: 136,000

— VLAN IDs: 4096

— Number of link aggregation groups (LAGs): 128

— Number of ports per LAG: 64

  • Number of FCoE VLANs/FC virtual fabrics: 4095

— Firewall filters:

  • Ingress: 768 Routed ACL (RACL), VLAN ACL (VACL), and Port ACL (PACL) rules

  • Egress: 1024 RACL, VACL, and PACL rules

— IPv4 unicast routes: 128,000 prefixes; 104,000 host routes

— IPv4 multicast routes: 52,000

— IPv6 multicast routes: 28,000

— IPv6 unicast routes: 98,000 prefixes; 52,000 host routes

— Address Resolution Protocol (ARP) entries: 32,000

— Generic routing encapsulation (GRE) tunnels: 1024

— MPLS labels: 16,000

— MPLS IPv4 L3 VPNs (Premium Services License): 2048

— Jumbo frame: 9216 bytes

— Spanning Tree Protocol (STP)

  • Multiple Spanning Tree Protocol (MSTP) instances: 64

  • VLAN Spanning Tree Protocol (VSTP) instances: 253

— Traffic mirroring

  • Mirroring destination ports per switch: 4

  • Maximum number of mirroring sessions: 4

  • Mirroring destination VLANs per switch: 4

Layer 2 features

— STP—IEEE 802.1D (802.1D-2004)

— Rapid Spanning Tree Protocol (RSTP) (IEEE 802.1w); MSTP (IEEE 802.1s)

— Bridge protocol data unit (BPDU) protect

— Loop protect

— Root protect

— RSTP and VSTP running concurrently

— VLAN—IEEE 802.1Q VLAN trunking

— Routed VLAN Interface (RVI)

— Port-based VLAN

— MAC address filtering

— QinQ

— VLAN translation

— Static MAC address assignment for interface

— Per VLAN MAC learning (limit)

  • MAC learning disable

  • Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)

  • IEEE 802.1AB Link Layer Discovery Protocol (LLDP)

Link aggregation

— Multichassis link aggregation (MC-LAG)

— Redundant trunk group (RTG)

— LAG load sharing algorithm—bridged or routed (unicast or multicast) traffic:

  • IP: Session Initiation Protocol (SIP), Dynamic Internet Protocol (DIP), TCP/UDP source port, TCP/UDP destination port

  • L2 and non-IP: MAC SA, MAC DA, Ethertype, VLAN ID, source port

  • FCoE packet: Source ID (SID), destination ID (DID), originator exchange ID (OXID), source port

Layer 3 features

— IStatic routing

— RIP v1/v2

— OSPF v1/v2

— OSPF v3

— Filter-based forwarding

— Virtual Router Redundancy Protocol (VRRP)

— IPv6

— Virtual routers

— Unicast RPF (uRPF)

— Loop-free alternate (LFA)

— BGP (Advanced Services or Premium Services license)

— IS-IS (Advanced Services or Premium Services license)

— Dynamic Host Configuration Protocol (DHCP) v4/v6 relay

— VR-aware DHCP

— IPv4/IPv6 over GRE tunnels (interface-based with decap/ encap and firewall-based with decap only)

Multicast

— Internet Group Management Protocol (IGMP) v1/v2

— Multicast Listener Discovery (MLD) v1/v2

— IGMP proxy, querier

— IGMP snooping

— MLD snooping

— Protocol Independent Multicast PIM-SM, PIM-SSM, PIM-DM

— Multicast Source Discovery Protocol (MSDP)

Security and filters

— Secure interface login and password

— RADIUS

— TACACS+

— Ingress and egress filters: Allow and deny, port filters, VLAN filters, and routed filters, including management port filters

— Filter actions: Logging, system logging, reject, mirror to an interface, counters, assign forwarding class, permit, drop, police, mark

— SSH v1, v2

— Static ARP support

— Storm control, port error disable, and autorecovery

— Control plane denial-of-service (DoS) protection

— DHCP snooping

Quality of Service (QoS)

— L2 and L3 QoS: Classification, rewrite, queuing

— Rate limiting:

  • Ingress policing: 1 rate 2 color, 2 rate 3 color

  • Egress policing: Policer, policer mark down action

  • Egress shaping: Per queue

— 10 hardware queues per port (8 unicast and 2 multicast)

— Strict priority queuing (LLQ), shaped-deficit weighted round-robin (SDWRR), weighted random early detection (WRED)

— 802.1p remarking

— Layer 2 classification criteria: Interface, MAC address, Ethertype, 802.1p, VLAN

— Congestion avoidance capabilities: WRED

— Trust IEEE 802.1p (ingress)

— Remarking of bridged packets

MPLS (Premium services license)

— Static label-switched paths (LSPs)

— RSVP-based signaling of LSPs

— LDP-based signaling of LSPs

— LDP tunneling (LDP over RSVP)

— MPLS class of service (CoS)

— MPLS access control list (ACL)/policers

— MPLS LSR support

— IPv4 L3 VPN (RFC 2547, 4364)

— MPLS fast reroute (FRR) - 1:1 Protection

Data Center Bridging (DCB)

— Priority-based flow control (PFC)—IEEE 802.1Qbb

— Data Center Bridging Exchange Protocol (DCBX), DCBx FCoE, and iSCSI type, length, and value (TLVs)

Fibre Channel over Ethernet (FCoE)

— FCoE transit switch (FIP snooping ACL installation)

— FCoE session path learning

— FCoE session health monitoring

— Graceful restart for FIP snooping

— FC-BB-6 VN2VN snooping

High availability

— Sub-second Bidirectional Forwarding Detection (BFD)

— Uplink failure detection (UFD)

Visibility and analytics

— Switched Port Analyzer (SPAN)

— Remote SPAN (RSPAN)

— Encapsulated Remote SPAN (ERSPAN)

— sFlow® v5

— Junos Telemetry Interface

Management and operations

— Role-based CLI management and access

— CLI via console, telnet, or SSH

— Extended ping and traceroute

— Junos OS configuration rescue and rollback

— Image rollback

— SNMP v1/v2/v3

— Junos XML management protocol

— High frequency statistics collection

— Beacon LED for port and system

— Automation and orchestration

— Zero touch provisioning (ZTP)

— OpenStack Neutron Plug-in

— Puppet

— Chef

— Python

— Junos OS event, commit, and OP scripts

Standards compliance

IEEE standards

— IEEE 802.1D

— IEEE 802.1w

— IEEE 802.1

— IEEE 802.1Q

— IEEE 802.1p

— IEEE 802.1ad

— IEEE 802.3ad

— IEEE 802.1AB

— IEEE 802.3x

— IEEE 802.1Qbb

— IEEE 802.1Qaz

T11 standards

— INCITS T11 FC-BB-5

Supported RFCs

— RFC 768 UDP

— RFC 783 Trivial File Transfer Protocol (TFTP)

— RFC 791 IP

— RFC 792 ICMP

— RFC 793 TCP

— RFC 826 ARP

— RFC 854 Telnet client and server

— RFC 894 IP over Ethernet

— RFC 903 RARP

— RFC 906 TFTP Bootstrap

— RFC 951 1542 BootP

— RFC 1058 Routing Information Protocol

— RFC 1112 IGMP v1

— RFC 1122 Host requirements

— RFC 1142 OSI IS-IS Intra-domain Routing Protocol

— RFC 1256 IPv4 ICMP Router Discovery (IRDP)

— RFC 1492 TACACS+

— RFC 1519 Classless Interdomain Routing (CIDR)

— RFC 1587 OSPF not-so-stubby area (NSSA) Option

— RFC 1591 Domain Name System (DNS)

— RFC 1745 BGP4/IDRP for IP—OSPF Interaction

— RFC 1772 Application of the Border Gateway Protocol in the Internet

— RFC 1812 Requirements for IP Version 4 routers

— RFC 1997 BGP Communities Attribute

— RFC 2030 SNTP, Simple Network Time Protocol

— RFC 2068 HTTP server

— RFC 2131 BOOTP/DHCP relay agent and Dynamic Host

— RFC 2138 RADIUS Authentication

— RFC 2139 RADIUS Accounting

— RFC 2154 OSPF w/Digital Signatures (Password, MD-5)

— RFC 2236 IGMP v2

— RFC 2267 Network ingress filtering

— RFC 2328 OSPF v2 (edge mode)

— RFC 2338 VRRP

— RFC 2362 PIM-SM (edge mode)

— RFC 2370 OSPF Opaque link-state advertisement (LSA) Option

— RFC 2385 Protection of BGP Sessions via the TCP Message Digest 5 (MD5) Signature Option

— RFC 2439 BGP Route Flap Damping

— RFC 2453 RIP v2

— RFC 2474 Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers

— RFC 2597 Assured Forwarding PHB (per-hop behavior) Group

— RFC 2598 An Expedited Forwarding PHB

— RFC 2697 A Single Rate Three Color Marker

— RFC 2698 A Two Rate Three Color Marker

— RFC 2796 BGP Route Reflection—An Alternative to Full Mesh IBGP

— RFC 2918 Route Refresh Capability for BGP-4

— RFC 3065 Autonomous System Confederations for BGP

— RFC 3376 IGMP v3 (source-specific multicast include mode only)

— RFC 3392 Capabilities Advertisement with BGP-4

— RFC 3446, Anycast RP

— RFC 3569 SSM

— RFC 3618 MSDP

— RFC 3623 Graceful OSPF Restart

— RFC 4271 Border Gateway Protocol 4 (BGP-4)

— RFC 4360 BGP Extended Communities Attribute

— RFC 4456 BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)

— RFC 4486 Subcodes for BGP Cease Notification Message

— RFC 4724 Graceful Restart Mechanism for BGP

— RFC 4812 OSPF Restart Signaling

— RFC 4893 BGP Support for Four-octet AS Number Space

— RFC 5176 Dynamic Authorization Extensions to RADIUS

— RFC 5396 Textual Representation of Autonomous System (AS) Numbers

— RFC 5668 4-Octet AS Specific BGP Extended Community

— RFC 5880 Bidirectional Forwarding Detection (BFD)

— Configuration Protocol (DHCP) server

Supported MIBs

— RFC 155 SMI

— RFC 1157 SNMPv1

— RFC 1212, RFC 1213, RFC 1215 MIB-II, Ethernet-Like MIB

— and TRAPs

— RFC 1850 OSPFv2 MIB

— RFC 1901 Introduction to Community-based SNMPv2

— RFC 2011 SNMPv2 for Internet protocol using SMIv2

— RFC 2012 SNMPv2 for transmission control protocol using SMIv2

— RFC 2013 SNMPv2 for user datagram protocol using SMIv2

— RFC 2233, The Interfaces Group MIB using SMIv2

— RFC 2287 System Application Packages MIB

— RFC 2570 Introduction to Version 3 of the Internet-standard Network Management Framework

— RFC 2571 An Architecture for describing SNMP Management Frameworks (read-only access)

— RFC 2572 Message Processing and Dispatching for the SNMP (read-only access)

— RFC 2576 Coexistence between SNMP Version 1, Version 2, and Version 3

— RFC 2578 SNMP Structure of Management Information MIB

— RFC 2579 SNMP Textual Conventions for SMIv2

— RFC 2580 Conformance Statements for SMIv2

— RFC 2665 Ethernet-like interface MIB

— RFC 2787 VRRP MIB

— RFC 2790 Host Resources MIB

— RFC 2819 RMON MIB

— RFC 2863 Interface Group MIB

— RFC 2932 IPv4 Multicast MIB

— RFC 3410 Introduction and Applicability Statements for Internet Standard Management Framework

— RFC 3411 An architecture for describing SNMP Management Frameworks

— RFC 3412 Message Processing and Dispatching for the SNMP

— RFC 3413 Simple Network Management Protocol (SNMP)—all MIBs are supported except the Proxy MIB

— RFC 3414 User-based Security Model (USM) for SNMP

— RFC 3415 View-based Access Control Model (VACM) for the SNMP

— RFC 3416 Version 2 of the Protocol Operations for the SNMP

— RFC 3417 Transport Mappings for the SNMP

— RFC 3418 Management Information Base (MIB) for the SNMP

— RFC 3584 Coexistence between Version 1, Version 2, and Version 3 of the Internet Standard Network Management Framework

— RFC 3826 The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model

— RFC 4188 Definitions of Managed Objects for Bridges

— RFC 4318 Definitions of Managed Objects for Bridges with Rapid Spanning Tree Protocol

— RFC 4363b Q-Bridge VLAN MIB

Environmental ranges

Parameters	QFX5200-32C	QFX5200-48Y-AFO[EOL]	QFX5200-48Y-AFI[EOL]
Operating temperature	32° to 104° F (0° to 40° C)	32° to104° F (0° to 40° C)	32° to 95° F (0° to 35° C)
Storage temperature	–40° to 158° F (–40° to 70° C)	–40° to 158° F (–40° to 70° C)	–40° to 158° F (–40° to 70° C)
Operating altitude	Up to 10000 ft. (3048 m)	Up to 5000 ft. (1524 m)	Up to 5000 ft. (1524 m)
Relative humidity operating	5 to 90% (noncondensing)	5 to 93% (noncondensing)	5 to 93% (noncondensing)
Relative humidity nonoperating	5 to 95% (noncondensing)	5 to 93% (noncondensing)	5 to 93% (noncondensing)
Seismic	Designed to meet GR-63, Zone 4 earthquake requirements	Designed to meet GR-63, Zone 4 earthquake requirements	Designed to meet GR-63, Zone 4 earthquake requirements

Power consumption

Parameters	QFX5200-32C	QFX5200-48Y[EOL]
Maximum power draw	*312 W AC	430 W (AC, DC), 1467 BTU/hr
Typical power draw	**184 W AC	382 W (AC, DC), 1303 BTU/hr

* Max power consumption is measured at 40°C ambient temperature with SR optics, and at 100% load with IMIX traffic.

** Typical power consumption is measured at 25°C ambient temperature with DACs, and at 50% load with IMIX traffic, excluding transceiver
Power consumption is subject to operating conditions and unit-to-unit variations.

Safety and compliance

Safety

— CAN/CSA-C22.2 No. 60950-1 Information Technology Equipment—Safety

— UL 60950-1 Information Technology Equipment—Safety

— EN 60950-1 Information Technology Equipment—Safety

— IEC 60950-1 Information Technology Equipment—Safety (All country deviations)

— EN 60825-1 Safety of Laser Products—Part 1: Equipment Classification

Electromagnetic compatibility

— 47 CFR Part 15, (FCC) Class A

— ICES-003 Class A

— EN 55022/EN 55032, Class A

— CISPR 22/CISPR 32, Class A

— EN 55024

— CISPR 24

— EN 300 386

— VCCI Class A

— AS/NZS CISPR 32, Class A

— KN32/KN35

— BSMI CNS 13438, Class A

— EN 61000-3-2

— EN 61000-3-3

— ETSI

— ETSI EN 300 019: Environmental Conditions & Environmental Tests for Telecommunications Equipment

— ETSI EN 300 019-2-1 (2000)—Storage

— ETSI EN 300 019-2-2 (1999)—Transportation

— ETSI EN 300 019-2-3 (2003)—Stationary Use at Weather-protected Locations

— ETSI EN 300 019-2-4 (2003)—Stationary Use at Non-Weather protected Locations

— ETS 300753 (1997)—Acoustic noise emitted by telecommunications equipment

Environmental compliance

— Restriction of Hazardous Substances (ROHS) 6/6

— China Restriction of Hazardous Substances (ROHS)

— Registration, Evaluation, Authorization and Restriction of Chemicals (REACH)

— Waste Electronics and Electrical Equipment (WEEE)

— Recycled material

— 80 Plus Silver PSU Efficiency

Telcot

— Common Language Equipment Identifier (CLEI) code

Ordering information

Product number	Description
Switch hardware	
QFX5200-32C-AFI2	QFX5200 (hardware only; software services sold separately), 32 QSFP+/QSFP28 ports, redundant fans, 2 AC power supplies, back-to-front airflow
QFX5200-32C-AFO2	QFX5200 (hardware only; software services sold separately), 32 QSFP+/QSFP28 ports, redundant fans, 2 AC power supplies, front-to-back airflow
QFX5200-32C-DC-AFI2	QFX5200 (hardware only; software services sold separately), 32 QSFP+/QSFP28 ports, redundant fans, 2 DC power supplies, back-to-front airflow
QFX5200-32C-DC-AFO2	QFX5200 (hardware only; software services sold separately), 32 QSFP+/QSFP28 ports, redundant fans, 2 DC power supplies, front-to-back airflow
JPSU-850W-AC-AFI	Juniper 850 W AC power supply for QFX5100-96S and QFX5200-32C (power supply-side airflow intake)
JPSU-850W-AC-AFO	Juniper 850 W AC power supply for QFX5100-96S and QFX5200-32C (power supply-side airflow exhaust)
JPSU-850W-DC-AFI	DC 850 W power supply, back-to-front airflow for QFX5100-96S and QFX5200-32C
JPSU-850W-DC-AFO	DC 850 W power supply, front-to-back airflow for QFX5100-96S and QFX5200-32C
QFX5200-32C-FANAFI	QFX5200-32C fan module, back-to-front airflow
QFX5200-32C-FANAFO	QFX5200-32C fan module, front-to-back airflow
QFX5200-48Y-AFI	QFX5200-48Y (base software services included), 48 SFP+/ SFP28 ports, redundant fans, 2 AC power supplies, back-to-front airflow
QFX5200-48Y-AFO	QFX5200-48Y (base software services included), 48 SFP+/ SFP28 ports, redundant fans, 2 AC power supplies, front-to-back airflow

Ordering information (continued)

Product number	Description
QFX5200-48Y-DC-AFI	QFX5200-48Y (base software services included), 48 SFP+/ SFP28 ports, redundant fans, 2 DC power supplies, back-to-front airflow
QFX5200-48Y-DC-AFO	QFX5200-48Y (base software services included), 48 SFP+/ SFP28 ports, redundant fans, 2 DC power supplies, front-to-back airflow
QFX520048Y-APSU-AO	Juniper 650 W AC power supply for QFX5200-48Y (power supply-side airflow intake)
QFX520048Y-APSU-AI	Juniper 650 W AC power supply for QFX5200-48Y (power supply-side airflow exhaust)
QFX520048Y-DPSU-AO	DC 650 W power supply, front-to-back airflow for QFX5200-48Y
QFX520048Y-DPSU-AI	DC 650 W power supply, back-to-front airflow for QFX5200-48Y
QFX520048Y-FAN-AO	QFX5200-48Y fan module, front-to-back airflow
QFX520048Y-FAN-AI	QFX5200-32C fan module, back-to-front airflow
EX-4PST-RMK	Adjustable 4-post rack-mount kit for EX Series, QFX5100, and QFX5200
QFX520048Y-RKMT	4 post rack-mount for QFX5200-48Y

Optics and transceivers

Up-to-date information on supported optics can be found on the Hardware Compatibility Tool at apps.juniper.net/hct/ product/?prd=QFX5200-32C. The QFX5200-32C line of Switches supports varying port speeds at 100GbE, 40GbE, and 10GbE with different transceiver options of direct attach copper cables, active optical cables, and breakout cables (DACBO and AOCBO).

Product number	Description
Software licenses	
S-QFX5K-C2-A1-3	Flex Sub Software, Class 2 QFX5000 line, Adv 1, Term: 3 Yrs
S-QFX5K-C2-A1-5	Flex Sub Software, Class 2 QFX5000 line, Adv 1, Term: 5 Yrs
S-QFX5K-C2-A1-P	Flex Software, Class 2 QFX5000 line, Adv 1, Perpetual
S-QFX5K-C2-A2-3	Flex Software, Class 2 QFX5000 line, Adv 2, Term: 3 Yrs
S-QFX5K-C2-A2-5	Flex Software, Class 2 QFX5000 line, Adv 2, Term: 5 Yrs
S-QFX5K-C2-A2-P	Flex Software, Class 2 QFX5000 line, Adv 2, Perpetual
S-QFX5K-C2-P1-3	Flex Software, Class 2 QFX5000 line, Prem 1, Term: 3 Yrs
S-QFX5K-C2-P1-5	Flex Software, Class 2 QFX5000 line, Prem 1, Term: 5 Yrs
S-QFX5K-C2-P1-P	Flex Software, Class 2 QFX5000 line, Prem 1, Perpetual

About HPE

HPE is a leader in essential enterprise technology, bringing together the power of AI, cloud, and networking to help organizations achieve more. As pioneers of possibility, our innovation and expertise advance the way people live and work. We empower our customers across industries to optimize operational performance, transform data into foresight, and maximize their impact. Unlock your boldest ambitions, with HPE. Discover more at HPE.com.

Chat now

Visit HPE.com

© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel is a trademark of Intel Corporation or its subsidiaries in the U.S. and/or other countries. sFlow is a registered trademark of InMon Corp. All third-party marks are property of their respective owners.

aa00150814ENW, Rev. 1

HEWLETT PACKARD ENTERPRISE

hpe.com