



QFX5220 SWITCH

Product Overview

Cloud providers and network operators are increasingly deploying scale-out, spine-and-leaf IP fabric architectures built on fixed-configuration switches to support growing east-west traffic in the data center.

The QFX5220 Switch is optimally suited for these high-speed, high-density, spine-and-leaf IP fabrics. Supporting 400GbE, 200GbE, 100GbE, 50GbE*, 40GbE, 25GbE, and 10GbE connections and offering an advanced L2, L3, and MPLS feature set, the QFX5220 enables cloud service providers and network operators to build large, next-generation IP fabrics that support network virtualization and intelligent traffic forwarding based on proven, Internet-scale technology.*

Product Description

The Juniper Networks® QFX5220 Switch is a next-generation, fixed-configuration spine-and-leaf switch. It offers flexible, cost-effective, high-density 400GbE, 200GbE*, 100GbE, 50GbE*, 40GbE, 25GbE, and 10GbE interfaces for server and intra-fabric connectivity, and delivers a versatile, future-proofed solution for today's data centers.

QFX5220 switches support advanced Layer 2, Layer 3, and MPLS features. For large public cloud providers—early adopters of high-performance servers to meet explosive workload growth—the QFX5220 supports very large, dense, and fast 400GbE 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 QFX5220 switch also provides a high radix-native 100GbE lean-spine option at reduced power and a smaller footprint.

Two QFX5220 models are available, supporting different configurations and use cases. Delivering 25.6 Tbps of bandwidth, both models are optimally designed for spine-and-leaf deployments in enterprise, HPC, service provider, and cloud data centers.

QFX5220-32CD: The QFX5220-32CD offers 32 ports in a low-profile 1 U form factor. High-speed ports support a wide variety of port configurations, including 400GbE, 200GbE, 100GbE, 25GbE, 40GbE, and 10GbE. The QFX5220-32CD is equipped with two AC or DC power supplies, providing 1+1 redundancy when all power supplies are present, and six hot-swappable fans offering ports-to-FRUs (AFO) or FRUs-to-ports (AFI) airflow options, providing (5x2+1)+1 redundancy.

QFX5220-128C: The QFX5220-128C offers 128 ports in a 4 U form factor. The high-speed ports support a wide variety of configurations, including 100GbE and 40GbE. The switch is equipped with four AC or DC power supplies, providing 2+2 redundancy when all power supplies are present, and six hot-swappable ports-to-FRUs (AFO) airflow fans, providing (5x2+1) +1 redundancy.

Both QFX5220 switch models include an Intel XeonD-1500 processor to drive the control plane, which runs the Juniper Networks Junos® OS Evolved operating system software.

*Reserved for future release

Product Highlights

The QFX5220 includes the following capabilities. Please refer to the Specifications section for currently shipping features.

Native 400GbE Configuration

The QFX5220-32CD offers 32 ports in a 1 U form factor. The high-speed ports support a wide variety of configurations, including 100GbE and 400GbE.

High-Density Configurations

The QFX5220 switches are optimized for high-density fabric deployments. The QFX5220-32CD provides an option of either 32 ports of 400GbE, 100GbE, or 40GbE, while the QFX5220-128C provides an option of either 128 QSFP28 100GbE ports or 64 40GbE QSFP ports.

Flexible Connectivity Options

The QFX5220 offers a choice of interface speeds for server and intra-fabric connectivity, providing deployment versatility and investment protection.

- **Rich automation capabilities:** The QFX5220 switches support a number of network automation features for plug-and-play operations, including zero-touch provisioning (ZTP), operations and event scripts, automatic rollback, and Python scripting.
- **Advanced Junos Evolved features:** The QFX5220 switch supports features such as BGP add-path, MPLS, L3 VPN, RoCEv2, and Multicast capabilities.

- **Junos Evolved software architecture:** The QFX5220 supports a modular Junos Evolved 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, support for seamless component upgrade without bringing the switch down, and support for containerization, enabling application deployment using LXC or Docker.

Deployment Options

The QFX5220-32CD can be deployed as a universal device in cloud data centers to support 100GbE server access and 400GbE spine-and-leaf configurations, optimizing data center operations by using a single device across multiple layers of the network (see Figure 1).

The QFX5220-128C is a high-radix 100GbE lean-spine switch optimized to aggregate 10GbE and 25GbE top-of-rack switches in these environments. Many cloud, service provider, data center, and enterprise networks are deploying 100GbE to handle growing demand. Figure 2 and Figure 3 show multiple use cases with the QFX5220-128C as a lean spine.

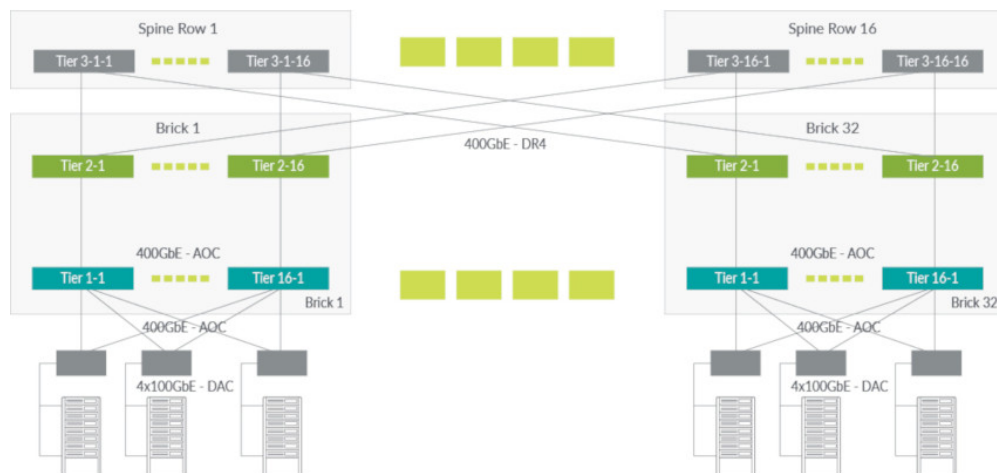


Figure 1: Typical cloud data center deployment for the QFX5220-32CD

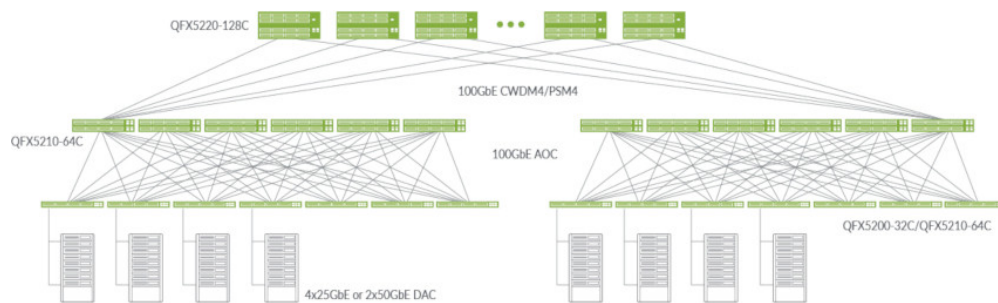


Figure 2: 100GbE fabric in a typical cloud data center

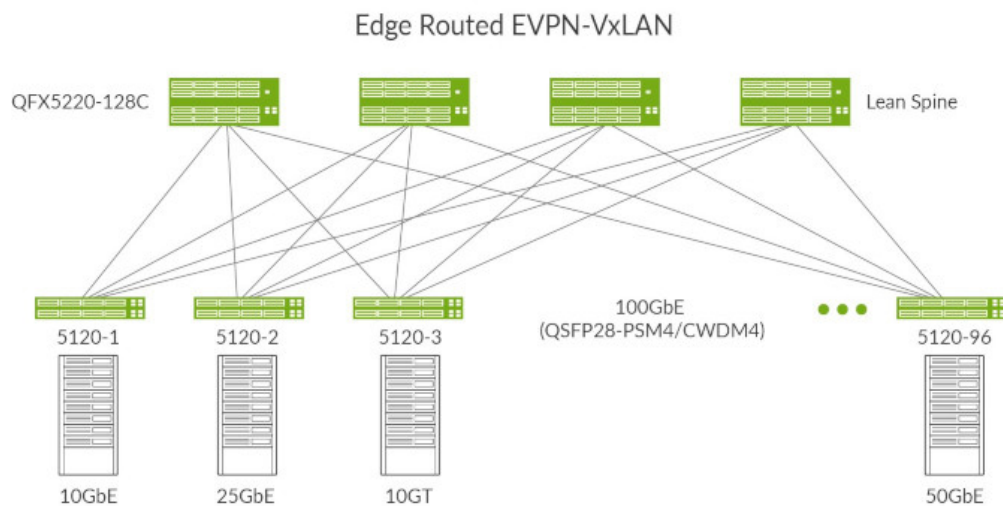


Figure 3: Private cloud data center with the QFX5220-128C as lean spine

Architecture and Key Components

The QFX5220 can be used in L3 fabrics and L2 networks. You can choose the architecture that best suits your deployment needs and easily adapt and evolve as requirements change over time. The QFX5220 serves as the universal building block for these two switching architectures, enabling data center operators to build cloud networks in their own way.

- Layer 3 fabric:** For customers looking to build scale-out data centers, a Layer 3 spine-and-leaf Clos fabric provides predictable, nonblocking performance and scale characteristics. A two-tier fabric built with QFX5220 switches as leaf devices and Juniper Networks QFX10000 line of Switches as the spine can scale to support up to 128 40GbE ports or 128 25GbE and/or 10GbE server ports in a single fabric. One of the most complicated tasks when building an IP fabric is assigning all the

implementation details, including IP addresses, BGP autonomous system numbers, routing policies, loopback address assignments, and others. Automating the creation of an IP fabric at a large scale is equally difficult. To address these challenges, Juniper has created the OpenClos project to provide free, 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 on GitHub, OpenClos takes a set of inputs that describe the shape and size of a data center and produces switch configuration files and a cabling plan

Junos Evolved ensures a high feature and bug fix velocity and provides first-class access to system state, allowing customers to run DevOps tools, containerized applications, management agents, specialized telemetry agents, and more.

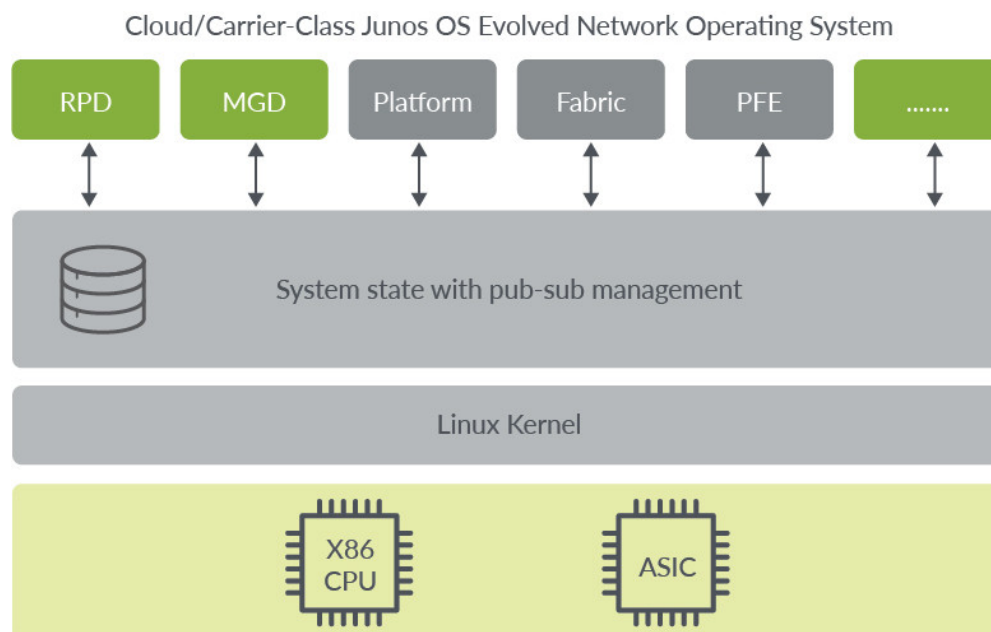


Figure 4: Cloud/Carrier-Class Junos OS Evolved Network Operating System

Management, Monitoring, and Analytics

Data Center Fabric Management: 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.

Apstra key features are:

- Automated deployment and zero-touch deployment
- Continuous fabric validation
- Fabric life-cycle management
- Troubleshooting using advanced telemetry

For more information on Apstra, see [Juniper Apstra](#).

Features and Benefits

- **Automation and programmability:** The QFX5220 supports numerous network automation features, including operations and event scripts and ZTP.
- **Cloud-level scale and performance:** The QFX5220 supports best-in-class cloud-scale L2/L3 deployments with a low latency of 750 ns and a superior scale and performance. This includes L2 support for 8192 media access control (MAC) addresses and Address Resolution Protocol (ARP) learning, which scales up to 32,000 entries at 500 frames per second. It also includes L3 support for 336,000 longest prefix match (LPM) routes and 380,000 host routes on IPv4. Additionally, the QFX5220 supports 130,000 LPM routes and 130,000 host routes on IPv6, 128-way equal-cost multipath (ECMP)

routes, and a filter that supports 768 (ingress) and 2558[†] (egress) exact match filtering rules. The QFX5220 supports up to 128 link aggregation groups, 4096 VLANs, and Jumbo frames of 9216 bytes. Junos Evolved provides configurable options through a CLI, enabling each QFX5220 to be optimized for different deployment scenarios.

[†] Ingress/egress scale numbers may be lower when used concurrently.

- **MPLS:** The QFX5220 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 service-level agreements (SLAs) at scale. The QFX5220 can also be deployed as a low-latency MPLS label-switching router (LSR) or MPLS provider edge (PE) router in smaller scale environments. The QFX5220, along with Juniper Networks QFX5100 and QFX5200 switches, are the most compact, low-latency, high-density, low-power family of switches to offer an MPLS feature set in the industry.
- **IEEE 1588 PTP Boundary Clock with Hardware Timestamping*:** IEEE 1588 PTP Transparent/Boundary Clock is supported on QFX5220, enabling accurate and precise sub-microsecond timing information in today's data center networks. In addition, the QFX5220 supports hardware timestamping; timestamps in Precision Time Protocol (PTP) packets are captured and inserted by an onboard field-programmable gate array (FPGA) on the switch at the physical (PHY) level.
- **Data packet timestamping*:** When the optional data packet timestamping feature is enabled, selected packets flowing through QFX5220 switches are timestamped with references

to the recovered PTP clock. When these packets are received by nodes in the network, the packet timestamping information can be mirrored onto monitoring tools for detailed analysis, helping identify bottlenecks in the network that cause latency. This information also helps with network performance analysis and record keeping for legal and compliance purposes, which is required by certain business transactions such as financial trading, video streaming, and research establishments.

- **RoCEv2:** As a switch capable of transporting data as well as storage traffic over Ethernet, the QFX5220 provides an IEEE data center bridging (DCB) converged network between servers with disaggregated flash storage arrays or an NVMe-enabled storage area network (SAN). The QFX5220 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. The RDMA over Converged Ethernet version 2 (RoCEv2) transit switch functionality, including priority-based flow control (PFC) and Data Center Bridging Capability Exchange (DCBX), are included as part of the default software.
- **Junos OS Evolved:** Junos Evolved is a native Linux operating system that incorporates a modular design of independent functional components and enables individual components to be upgraded independently while the system remains operational. Component failures are localized to the specific component involved and can be corrected by upgrading and restarting that specific component without having to bring down the entire device.
- **Retained state:** State is the retained information or status pertaining to physical and logical entities. It includes both

operational and configuration state, comprising committed configuration, interface state, routes, hardware state, and what is held in a central database called the distributed data store (DDS). State information remains persistent, is shared across the system, and is supplied during restarts.

- **Feature support:** All key networking functions such as routing, bridging, management software, and management plane interfaces, as well as APIs such as CLI, NETCONF, Juniper Extension Toolkit (JET), Junos Telemetry Interface (JTI), and the underlying data models, resemble those supported by Junos. This ensures compatibility and eases the transition to Junos Evolved.

Junos Telemetry Interface

The QFX5220 supports Junos Telemetry Interface (JTI), a modern telemetry streaming tool that provides performance monitoring in complex, dynamic data centers. Streaming data to a performance management system lets network administrators measure trends in link and node utilization and troubleshoot issues such as network congestion in real time.

JTI provides:

- Application visibility and performance management by provisioning sensors to collect and stream data and analyze the 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.



QFX5220-128C



QFX5220-32CD

Specifications

Hardware

Table 1: QFX5220 System Capacity

Specification	QFX5220-32CD	QFX5220-128C
System throughput	Up to 25.6 Tbps (bidirectional)	Up to 25.6 Tbps (bidirectional)
Forwarding capacity	8 billion packets per second	8 billion packets per second
Port density	32 ports of QSFP56-DD 400GbE	128 ports of QSFP28 100GbE
SFP+/SFP28	2 SFP+ transceiver ports for in-band network management	2 SFP+ transceiver ports for in-band network management

Table 2: QFX5220 System Specifications

Specification	QFX5220-32CD	QFX5220-128C
Dimensions (W x H x D)	17.26 x 1.72 x 21.1 in. (43.8 x 4.3 x 53.59 cm)	17.26 x 6.88 x 29 in. (43.8 x 17.47 x 73.66 cm)
Rack units	1 U	4 U
Weight	24.5 lb (11.11 kg) with power supplies and fans installed	98 lb (44.44 kg) with 4 power supplies and 6 fan trays installed
Operating system	Junos OS Evolved	Junos OS Evolved
CPU	Intel Xeon D-1518	Intel Xeon D-1518
Power	<ul style="list-style-type: none"> Redundant (1+1) hot-pluggable 1600 W AC/DC power supplies (2n) 115-240 V single phase AC power -48 to -60 V DC power 	<ul style="list-style-type: none"> Redundant (1+1) hot-pluggable 1600 W AC/DC power supplies 115-240 V single phase AC power -48 to -60 V DC power
Cooling	<ul style="list-style-type: none"> Ports-to-FRUs (AFO) and FRUs-to-ports (AFI) cooling Redundant (5x2+1)+1 hot-pluggable fan modules with variable speed to minimize power draw 	<ul style="list-style-type: none"> Ports-to-FRUs (AFO) cooling Redundant (5x2+1) + 1 hot-pluggable fan modules with variable speed to minimize power draw
Total packet buffer	64 MB	64 MB
Recommended Software Version	Junos OS Evolved 19.2R1 and Later	Junos OS Evolved 19.3R1 and Later
Warranty	Juniper standard one-year warranty	Juniper standard one-year warranty

Software

- MAC addresses per system: 8192
- VLAN IDs: 3968 (QFX5220-32CD) 3952 (QFX5220-128C)
- Number of link aggregation groups (LAGs): 128
- Number of ports per LAG: 64
- Firewall filters:
 - Ingress: 768 Routed ACL (RACL), VLAN ACL (VAACL), and Port ACL (PAACL) rules
 - Egress: 2558[†] RAACL; 512 VAACL and PAACL rules
- IPv4 unicast routes: 380,000 prefixes; 380,000 host routes
- IPv6 unicast routes: 130,000 prefixes; 130,000 host routes
- Address Resolution Protocol (ARP) entries: 32,000
- Generic routing encapsulation (GRE) tunnels: 2000
- MPLS labels: 16,000
- Jumbo frame: 9216 bytes
- Traffic mirroring
 - Mirroring destination ports per switch: 4
 - Maximum number of mirroring sessions: 4
 - Mirroring destination VLANs per switch: 4

[†]Note: Ingress/Egress scale numbers may be lower when used concurrently.

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*
- Static MAC address assignment for interface
- MAC learning disable
- Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)

*Reserved for future release

Link Aggregation

- 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, Ether type, VLAN ID, source port

Layer 3 Features

- Static routing
- OSPF v1/v2
- OSPF v3
- Filter-based forwarding
- Virtual Router Redundancy Protocol (VRRP)*
- IPv6
- Virtual routers
- 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 only)

*Reserved for future release

Multicast*

- Internet Group Management Protocol (IGMP) v1/v2
- Multicast Listener Discovery (MLD) v1/v2
- IGMP proxy, querier
- IGMP v1/v2/v3 snooping
- Intersubnet multicast using IRB interface
- MLD snooping
- Protocol Independent Multicast PIM-SM, PIM-SSM, PIM-DM, PIM-Bidir*
- Multicast Source Discovery Protocol (MSDP)*

*Reserved for future release

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, loopback filters for control plane protection
- 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

*Reserved for future release

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, per port
- 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, ECN
- Trust IEEE 802.1p (ingress)
- Remarking of bridged packets
- Configurable shared buffer and buffer monitoring

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)

*Reserved for future release

Data Center Bridging (DCB)*

- Priority-based flow control (PFC)—IEEE 802.1Qbb
- Data Center Bridging Exchange Protocol (DCBX)*

*Reserved for future release

High Availability

- Bidirectional Forwarding Detection (BFD)
- Uplink failure detection (UFD)*

*Reserved for future release

Visibility and Analytics

- Switched Port Analyzer (SPAN)
- Remote SPAN (RSPAN)
- Encapsulated Remote SPAN (ERSPAN)
- sFlow v5
- Junos Telemetry Interface

Management and Operations

- Contrail Networking*
- Role-based CLI management and access
- CLI via console, telnet, or SSH
- Extended ping and traceroute
- Junos OS Evolved configuration rescue and rollback
- Image rollback
- SNMP v1/v2/v3
- Junos OS Evolved XML management protocol
- High frequency statistics collection
- Automation and orchestration
- Zero-touch provisioning (ZTP)
- Python
- Junos OS Evolved event, commit, and OP scripts

*Reserved for future release

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

*Reserved for future release

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 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 Source-specific multicast (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
- 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
- RFC 3414 User-based Security Model (USM) for SNMPv3
- 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

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)

Environmental Ranges

Parameters	QFX5220-32CD	QFX5220-128C
Operating temperature	32° to 104° F (0° to 40° C)	32° to 104° F (0° to 40° C)
Storage temperature	-40° through 158° F	-40° through 158° F
Operating altitude	Up to 6000 feet (1828.8 meters)	Up to 6000 feet (1828.8 meters)
Relative humidity operating	5 to 90% (noncondensing)	5 to 90% (noncondensing)
Relative humidity nonoperating	5 to 95% (noncondensing)	5 to 95% (noncondensing)
Seismic	Designed to meet GR-63, Zone 4 earthquake requirements	Designed to meet GR-63, Zone 4 earthquake requirements

Maximum Thermal Output

Parameters	QFX5220-32CD	QFX5220-128C
Maximum power draw	115-127 V: 973 W; 220-240 V: 958 W	115-127 V: 2023 W; 220-240 V: 1990 W
Typical power draw	115-127 V: 730 W; 220-240 V: 775 W	115-127 V: 1433 W; 220-240 V: 1394 W

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

Security

- FIPS/CC*
- TAA*

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 Weatherprotected Locations
- ETSI EN 300 019-2-4 (2003)—Stationary Use at NonWeather-protected Locations
- ETS 300753 (1997)—Acoustic noise emitted by telecommunications equipment

Environmental Compliance



Restriction of Hazardous Substances (ROHS) 6/6



Silver PSU Efficiency



Recycled material



Waste Electronics and Electrical Equipment (WEEE)



Registration, Evaluation, Authorisation and Restriction of Chemicals (REACH)



China Restriction of Hazardous Substances (ROHS)

Telco

- Common Language Equipment Identifier (CLEI) code

Juniper Networks Services and Support

Juniper Networks is the leader in performance-enabling services that are designed to accelerate, extend, and optimize your high-performance network. Our services allow you to maximize operational efficiency while reducing costs and minimizing risk, achieving a faster time to value for your network. Juniper Networks ensures operational excellence by optimizing the network to maintain required levels of performance, reliability, and availability. For more details, please visit www.juniper.net/us/en/products-services.

Ordering Information

Product Number	Description
Hardware	
QFX5220-32CD-AFI	QFX5220 (hardware with base software), 32 QSFP-DD/QSFP+/QSFP28 ports, redundant fans, 2 AC power supplies, back-to-front airflow
QFX5220-32CD-AFO	QFX5220 (hardware with base software), 32 QSFP-DD/QSFP+/QSFP28 ports, redundant fans, 2 AC power supplies, front-to-back airflow
QFX5220-32CD-D-AFI	QFX5220 (hardware with base software), 32 QSFP-DD/QSFP+/QSFP28 ports, redundant fans, 2 DC power supplies, back-to-front airflow
QFX5220-32CD-D-AFO	QFX5220 (hardware with base software), 32 QSFP-DD/QSFP+/QSFP28 ports, redundant fans, 2 DC power supplies, front-to-back airflow
QFX5220-128C-AFO	QFX5220 (hardware with base software), 128 QSFP+/QSFP28 ports, redundant fans, 4 AC power supplies, front-to-back airflow

Product Number	Description
QFX5220-128C-D-AFO	QFX5220 (hardware with base software), 128 QSFP+/QSFP28 ports, redundant fans, 4 DC power supplies, front-to-back airflow
JPSU-1600W-1UACAFI	QFX5220-32CD-AFI 1 U AC power supply unit
JPSU-1600W-1UACAFO	QFX5220-32CD-AFO 1 U AC power supply unit
JPSU-1600W-1UDCAFI	QFX5220-32CD-D-AFI 1 U DC power supply unit
JPSU-1600W-1UDCAFO	QFX5220-32CD-D-AFO 1 U DC power supply unit
QFX5220-32CD-4PRMK	4-Post Rack Mount Kit for QFX5220-32CD
RKMT-4PST-4U	4-Post Rack Mount Kit for QFX5220-128C
QFX5220-32CD-FANAI	Airflow in (AFI) back-to-front airflow fans for QFX5220-32CD
QFX5220-32CD-FANAO	Airflow out (AFO) front-to-back airflow fans for QFX5220-32CD
QFX5220-128C-FANAO	Airflow out (AFO) front-to-back airflow fans for QFX5220-128C
QFX5220-14I-EMI-DR	QFX5220-128C 14 in. EMI door

Optics and Transceivers

QDD-400G-DAC-1M	QSFP56-DD 400GbE DAC 1 M
QDD-400G-DAC-2P5M	QSFP56-DD 400GbE DAC 2.5 M
QDD-4x100G-FR	QSFP-DD 4x100GBASE-FR breakout 1310 nm PAM4 transceiver module, 2 km reach
QDD-400G-DR4	QSFP-DD 400GBASE-DR4 / 4x100GBASE-DR 1310 nm PAM4 transceiver module, 500 m reach
QDD-400G-FR4	QSFP-DD 400GBASE-FR4 1310 nm PAM4 transceiver module, 2 km reach
QDD-400G-AOC-XM	400GbE QSFP56-DD active optical cable of XM (X=1,3,5,7,10,15,20,30)
JNP-100G-4x25G-1M	100GbE QSFP28 to 4x25GbE SFP28 passive direct attach copper breakout cable, length: 1 m
JNP-100G-4x25G-3M	100GbE QSFP28 to 4x25G SFP28 passive direct attach copper breakout cable, length: 3 m
JNP-QSFP-100G-BXSR*	100GbE SR-bidirectional QSFP transceiver, LC, 100 M, OM4 multimode fiber-optic (MMF)

Software Licenses SKUs

S-QFX5K-C3-A1-X (X=3,5)	Base L3 Software Subscription (X Years; X=3,5) License for QFX5220-32CD/QFX5220-128C
S-QFX5K-C3-A2-X (X=3,5)	Advanced Software Subscription (X Years; X=3,5) License for QFX5220-32CD/QFX5220-128C
S-QFX5K-C3-P1-X (X=3,5)	Premium Software Subscription (X Years; X=3,5) License for QFX5220-32CD/QFX5220-128C

About Juniper Networks

At Juniper Networks, we are dedicated to dramatically simplifying network operations and driving superior experiences for end users. Our solutions deliver industry-leading insight, automation, security and AI to drive real business results. We believe that powering connections will bring us closer together while empowering us all to solve the world's greatest challenges of well-being, sustainability and equality.

Corporate and Sales Headquarters

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089 USA

Phone: 888.JUNIPER (888.586.4737)

or +1.408.745.2000

www.juniper.net

APAC and EMEA Headquarters

Juniper Networks International B.V. Boeing
Avenue 240 1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands

Phone: +31.207.125.700

