



QFX5210 SWITCH

Product Overview

The QFX5210 Switch is a flexible, 64-port, fixed-configuration switch designed for spine-and-leaf applications in next-generation IP fabric networks. Featuring line-rate, high-density QSFP+/QSFP28 ports supporting 10GbE, 25GbE, 40GbE, and 100GbE connections, the QFX5210 offers an advanced Layer 2, Layer 3, and MPLS feature set. This allows network operators to build large, high-density IP fabrics that support network virtualization and intelligent traffic forwarding based on proven, Internet-scale technology, while enabling easy upgrades to 25GbE servers.

Additionally, the QFX5210-64C-S platform includes support for SONiC, continuing Juniper's commitment to disaggregation, allowing customers to innovate quickly and easily while simplifying operations.

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® QFX5210 Switch is a next-generation, fixed-configuration spine/leaf switch that offers flexible, cost-effective, high-density 10GbE, 25GbE, 40GbE, and 100GbE interfaces for IP fabric networks, providing deployment versatility, investment protection, and future-proofing for today's data centers. The QFX5210 Switch's high port density also contributes to simplified Clos architectures, eliminating network tiers while reducing complexity and cost.

Like other Juniper Networks QFX Series Switches, the QFX5210 supports advanced L2, L3, and MPLS features. For large public cloud providers—among the first to adopt 25GbE servers to meet explosive workload growth—the QFX5210 enables very large, dense, and fast IP fabrics based on proven Internet scale technology.

The QFX5210 runs the same reliable, high-performance Junos® operating system used by network operators around the world.

QFX5210 Deployment Options

Table 1 shows some of the many QFX5210 deployment options, including top-of-rack access as well as spine-and-leaf configurations in a centralized or distributed gateway architecture.

Table 1. QFX5210 Deployment Options

Port Combinations	Deployment
64x100GbE	Spine/leaf
96x25GbE + 8x100GbE	25GbE access

The QFX5210 can be positioned as a top-of-rack device in a centralized gateway architecture (also known as centrally routed bridging, or CRB). Figure 1 shows the QFX5210 deployed as an access switch with the Juniper Networks QFX10000 line of Switches acting as the spine device configured as a central gateway. In this topology, the QFX10000 can also be configured as a collapsed spine and edge device, acting as a central gateway for Ethernet VPN (EVPN)-Virtual Extensible LAN (VXLAN), MPLS, and other tunneling protocols. The 100GbE ports on the QFX5210 can be channelized to support 4x25GbE or 4x10GbE downlinks.




Figure 1: QFX5210 and QFX10002 in a leaf-spine deployment with centralized gateway

The QFX5210 can also be positioned as a leaf spine device in a distributed gateway architecture (also known as edge-routed bridging, or ERB). Figure 2 shows QFX5120 switches deployed as access switches configured as distributed gateways and QFX5210 switches deployed as spine switches.




Figure 2: QFX5120 and QFX5210 leaf-spine deployment with distributed gateway

Architecture and Key Components

The QFX5210 is a compact, 2 U platform featuring 64 quad small form-factor pluggable plus (QSFP+) transceivers or QSFP28 ports, with two redundant 1100 W AC/DC power supplies and four fan trays.

A QSFP28 module can be used to provide a 1x100GbE connection or 4x25GbE connections using breakout cables. A QSFP+ module can be used to provide a 1x40GbE connection or 4x10GbE connections using breakout cables.

The QFX5210 features a high-throughput 12.8 Tbps Packet Forwarding Engine (PFE). Control plane performance is further enhanced with the Intel Broadwell-DE, 4 Core, 2.2 GHz CPU, and 16 GB DDR4, 100 GB SSD storage.

The QFX5210 can be used in L3 fabrics and L2 networks with multichassis link aggregation (MC-LAG). Customers can choose the architecture that best suits their deployment needs and easily adapt and evolve as requirements change over time. The QFX5210 switch serves as a universal building block for these switching architectures, enabling data center operators to build cloud networks their way.

- IP Fabric:** For customers looking to build scale-out data centers, an L3 spine-and-leaf Clos fabric is ideal due to its predictable, nonblocking performance and scale characteristics. For example, a two-tier fabric built with the QFX5210 Switch as leafs and 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.
- Overlays:** Customers can deploy overlay networks to provide L2 adjacencies for applications over L3 fabrics. The overlay networks use VXLAN in the data plane and EVPN to program the overlays. The overlays can operate without a controller or can be orchestrated with a fabric management orchestrator to provide L2/L3 overlay virtual networking and security for bare-metal servers and virtual workloads. QFX5210 switches can also integrate with VMware NSX.
- MC-LAG:** The QFX5210 supports the MC-LAG protocol between two switches, eliminating Spanning Tree Protocol (STP) in traditional L2 networks when deployed in the aggregation layer. MC-LAG's active/active operation ensures complete bandwidth utilization between the network's access and aggregation layers, while the dual control plane technology ensures the highest availability for applications.

Juniper HealthBot

Combining the power of telemetry, programmability, advanced algorithms, and machine learning, the Juniper HealthBot solution revolutionizes data analytics. Integrated with the Juniper Telemetry Interface (JTI), HealthBot aggregates large volumes of real-time telemetry data and correlates analytics to provide predictive insights that offer a multidimensional view across the entire network, as well as the applications running on it. Open programmability supports customized playbooks, enabling service providers and enterprises to build highly customized health monitoring and diagnostics workflows.

By focusing on actionable insights, simplified consumption, and a programmable framework supported by open-sourced data pipelines and collectors for data ingestion, HealthBot democratizes network analytics and encourages collaboration across business units, ultimately enhancing agility and innovation across the entire ecosystem.

Features and Benefits

- **Flexible connectivity options:** The QFX5210 offers a choice of 10GbE, 25GbE, 40GbE, and 100GbE interface speeds for server and intra-fabric connectivity, providing deployment versatility and investment protection.
- **Automation and programmability:** The QFX5210 supports numerous network automation features, including operations and event scripts, ZTP, and Juniper plug-ins for OpenStack Neutron.
- **Flexible Forwarding Table:** The QFX5210 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. Junos OS provides configurable options through a CLI, enabling each QFX5210 to be optimized for different deployment scenarios.
- **MPLS:** The QFX5210 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 QFX5210 can also be deployed as a low-latency MPLS label-switching router (LSR) or MPLS provider edge (PE) router in smaller scale environments. The QFX5210, along with the Juniper Networks QFX5220, QFX5200, QFX5120, QFX5110, and QFX5100 switches, is among the industry's most compact, low-latency, high-density, low-power family of switches to offer an MPLS feature set.

- **RoCEv2:** As a switch capable of transporting data as well as storage traffic over Ethernet, the QFX5210 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 QFX5210 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 DiffServ code point (DSCP) priority-based flow control (PFC) and Data Center Bridging Capability Exchange (DCBX) with Explicit Congestion Notification (ECN) are included as part of the default software.

QFX5210-64C-S Hardware with SONiC Support

The QFX5210-64C-S Switch supports Juniper's SONiC implementation, delivering best-of-breed hardware and routing for cloud providers while taking advantage of the flexibility, resiliency, and cost savings of SONiC's open and disaggregated architecture.

The SONiC-enabled QFX5210-64C-S, which plugs seamlessly into a unified SONiC network infrastructure, leverages the existing QFX5210-64C hardware and includes support for ONIE and IDEEPROM TLV formats. ONIE and SONiC images are installed on the hardware at the factory. The Juniper-provided SONiC image includes the platform and FRU device drivers only; it does not include Juniper software features such as L2, L3, MPLS, EVPN, VXLAN, etc. The platform and FRU device drivers are uploaded to the SONiC Github open-source repository. Customers can use their private SONiC patches by building and upgrading the ONIE and SONiC images on the switch.

Juniper complements the SONiC OS with the containerized routing protocol process (cRPD), a full-function routing and management stack packaged as a container, ensuring a common end-to-end routing experience across different tiers in the data center. In addition, cRPD enables high-performance telemetry, automation, and programmability in a lightweight deployment.



QFX5210

Specifications

Hardware

Description	Measurement
System throughput	Up to 12.8 Tbps (bidirectional)
Forwarding capacity	Up to 4.2 Bpps
QSFP+/QSFP28 ports	64 QSFP+ or QSFP28
25GbE (breakout cable, QSFP28)	128
10GbE (breakout cable, QSFP+)	128+2
Dimensions (W x H x D)	17.26 x 3.45 x 24.1 in. (43.84 x 8.77 x 61.2 cm) (depth including fan handle; otherwise 22.83 in. (58 cm))
Rack units	2 U
Weight	14.1 kg
Operating system	Junos OS
CPU	Intel Broadwell-DE, 4 Core, 2.2 GHz, 16 GB DDR4, 100 GB SSD
Management interfaces	1 RJ-45 management port provided
Power	Redundant (1+1) hot-pluggable 1100 W AC/DC power supplies 110-240 V single phase AC power -36 to -72 V DC power
Cooling	Front-to-back and back-to-front cooling Redundant (N+1) hot-pluggable fan modules with variable speed to minimize power draw
Total packet buffer	42 MB
Warranty	Juniper standard one-year warranty

Software (Non-SONiC Versions)

Performance Scale (Unidimensional)

- MAC addresses per system: 264,000
- VLAN IDs: 4096 (3 VLANs reserved for internal use)
- Number of link aggregation groups (LAGs): 64
- Number of ports per LAG: 64
- Firewall filters
 - Ingress: 768 routed ACL (RACL) rules, 768 VLAN ACL (VACL) rules, 768 port ACL (PACL) rules
 - Egress: 1024 RACL, VACL, and PACL rules
- IPv4 route capacity: 262,140
- IPv4 host table capacity: 204,750
- IPv6 route capacity: 172,016
- IPv6 host table capacity: 102,339
- Address Resolution Protocol (ARP) entries: 49,000
- Generic routing encapsulation (GRE) tunnels: 2040
- MPLS labels: 32,000
- MPLS IPv4 L3 VPNs (Advanced Services License): 2048
- Jumbo frame: 9216 bytes
- Spanning Tree Protocol (STP)
 - Multiple Spanning Tree Protocol (MSTP) instances: 64
 - VLAN Spanning Tree Protocol (VSTP) instances: 509

- Traffic mirroring
 - Mirroring destination ports per switch: 4
 - Maximum number of mirroring sessions: 4
 - Mirroring destination VLANs per switch: 4

Layer 2 Features

- EVPN-VXLAN
- 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
- Q-in-Q
- 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
 - Layer 2 and non-IP: MAC SA, MAC DA, Ethertype, VLAN ID, source port

Layer 3 Features

- Static 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)
- DSCP-based PFC

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-Bidir
- Multicast Source Discovery Protocol (MSDP)
- Selective Multicast Ethernet Tag (SMET)—EVPN Type-6
- EVPN—Type-7/8

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, 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
- Trust IEEE 802.1p (ingress)
- Remarking of bridged packets

MPLS (Advanced 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)

Data Center Bridging (DCB)

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

High Availability

- Sub-second Bidirectional Forwarding Detection (BFD)
- Uplink failure detection (UFD)

Visibility and Analytics

- Local port mirroring (comparable to SPAN)
- Remote port mirroring (comparable to RSPAN)
- Remote port mirroring with a GRE tunnel (comparable to ERSPAN)
- Cloud Analytics Engine flow path analysis
- sFlow v5

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

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
- FC 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 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

Environmental Ranges

Parameters	QFX5210
Operating temperature	32° to 104° F (0° to 40° C)
Storage temperature	-40° to 158° F (-40° to 70° C)
Operating altitude	AFO models: Up to 6000 ft. (1828 m) AFI models: Sea level only
Relative humidity operating	5 to 90% (noncondensing)
Relative humidity nonoperating	5 to 95% (noncondensing)
Seismic	Designed to meet GR-63, Zone 4 earthquake requirements

Maximum Thermal Output

Parameters	QFX5210
Maximum power draw	960 W (AC), 980 W (DC)
Typical power draw	680 W (AC), 680 W (DC)

Safety and Compliance

Safety

- CAN/CSA-C22.2 No. 60950-1
- UL 60950-1 (2nd Edition)
- IEC 60950-1: 2005/A2:2013

Electromagnetic Compatibility

- EN 300 386
- EN 55032/CISPR 32, Class A
- EN 55022/CISPR 22, Class A
- EN 55024/CISPR 24, Class A
- FCC 47 CFR Part 15, Class A
- ICES-003, Class A AS/NZS CISPR 32
- VCCI-CISPR 32, Class A
- BSMI CNS 13438
- KN32/KN35
- 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

Environmental Compliance



Restriction of Hazardous Substances (ROHS) 6/6



China Restriction of Hazardous Substances (ROHS)



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



Waste Electronics and Electrical Equipment (WEEE)



Recycled material



80 Plus Silver PSU Efficiency

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.

Support and services for the SONiC-enabled QFX5210-64C-S includes fixes for the platform and FRU device drivers. 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

For more information, please contact your Juniper Networks representative.

Model Number	Description
Switch Hardware	
QFX5210-64C-AFI	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, back-to-front airflow
QFX5210-64C-AFI2 (Use with Flex licensing)	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, back-to-front airflow, Flex Transform
QFX5210-64C-DC-AFI	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 DC power supplies, back-to-front airflow
QFX5210-64C-D-AFI2 (Use with Flex licensing)	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 DC power supplies, back-to-front airflow, Flex Transform
QFX5210-64C-AFO	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, front-to-back airflow
QFX5210-64C-AFO2 (Use with Flex licensing)	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, front-to-back airflow, Flex Transform
QFX5210-64C-DC-AFO	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 DC power supplies, front-to-back airflow
QFX5210-64C-D-AFO2 (Use with Flex licensing)	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 DC power supplies, front-to-back airflow, Flex Transform
QFX5210-64C-SAFI	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, back-to-front airflow, with SONiC
QFX5210-64C-SAFO	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, front-to-back airflow, with SONiC
QFX5210-64C-DCSAFI	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 DC power supplies, back-to-front airflow, with SONiC
QFX5210-64C-DCSAFO	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 DC power supplies, front-to-back airflow, with SONiC
QFX5210-64C-AFO-T	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, front-to-back airflow, TAA
QFX5210-64C-AFO-T2 (Use with Flex licensing)	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, front-to-back airflow, TAA, Flex Transform
QFX5210-64C-AFI-T	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, back-to-front airflow, TAA
QFX5210-64C-AFI-T2 (Use with Flex licensing)	QFX5210, 64 QSFP+/QSFP28 ports, reduced latency, redundant fans, 2 AC power supplies, back-to-front airflow, TAA, Flex Transform
JPSU-1100W-AC-AFI	PSU, 1100 W AC, 12 V output, FRU, back-to-front airflow
JPSU-1100W-AC-AFO	PSU, 1100 W AC, 12 V output, FRU, front-to-back airflow
QFX5210-FANAFI	QFX5210-FANAFI FAN model, back-to-front airflow
QFX5210-FANAFO	QFX5210-FANAFO FAN model, front-to-back airflow
EX4500-4PST-RMK	4-post rack mount kit

Model Number	Description
Optics and Transceivers	
QFX-QSFP-40G-SR4	QSFP+ 40GBASE-SR4 40 gigabit optics, 850 nm for up to 150 m transmission on multimode fiber-optic (MMF)
QFX-QSFP-DAC-3M	QSFP+-to-QSFP+ Ethernet DAC (twinax copper cable), 3 m, passive
QFX-QSFP-DAC-1M	QSFP+-to-QSFP+ Ethernet DAC (twinax copper cable), 1 m, passive
JNP-QSFP-100G-SR4	QSFP28 100GBASE-SR4 optics for up to 100 m transmission over parallel MMF
JNP-QSFP-100G-LR4	QSFP28 100GBASE-LR4 optics for up to 10 km transmission over serial SMF
JNP-100G-DAC-3M	QSFP28-to-QSFP28 Ethernet DAC (twinax copper cable), 3 m
JNP-100G-DAC-1M	QSFP28-to-QSFP28 Ethernet DAC (twinax copper cable), 1 m

Optics and Transceivers (non-SONiC versions only)

QFX-SFP-10GE-LR	SFP+ 10GBASE-LR 10GbE optics, 1310 nm for 10 km transmission on single-mode fiber-optic (SMF) (for management port)
JNP-QSFP-4X10GE-IR	QSFP+ 40GBASE optics, up to 1.4 km transmission on parallel single mode (4x10GbE long reach up to 1.4 km)
JNP-QSFP-4X10GE-LR	QSFP+ 40GBASE optics, up to 10 km transmission on parallel single mode (4x10GbE long reach up to 10 km)
QFX-QSFP-DACBO-3M	QSFP+-to-SFP+ 10GbE DAC breakout copper (twinax copper cable), 3 m
QFX-QSFP-DACBO-1M	QSFP+-to-SFP+ 10GbE DAC breakout copper (twinax copper cable), 1 m
JNP-QSFP-100G-CWDM	QSFP28 100GBASE-CWDM4 optics for up to 2 km transmission over serial SMF
JNP-QSFP-100G-PSM4	QSFP28 100GBASE-PSM4 optics for up to 500 transmission over serial SMF
JNP-100G-4X25G-1M	QSFP28-to-SFP+ 25 Gigabit Ethernet DAC breakout copper (twinax copper cable), 1 m
JNP-100G-4X25G-3M	QSFP28-to-SFP+ 25 Gigabit Ethernet DAC breakout copper (twinax copper cable), 3 m
JNP-100G-AOC-1M	100GbE QSFP-to-QSFP active optical cables, 1 m
JNP-100G-AOC-3M	100GbE QSFP-to-QSFP active optical cables, 3 m
JNP-100G-AOC-5M	100GbE QSFP-to-QSFP active optical cables, 5 m
JNP-100G-AOC-10M	100GbE QSFP-to-QSFP active optical cables, 10 m
JNP-100G-AOC-20M	100GbE QSFP-to-QSFP active optical cables, 20 m
JNP-100G-AOC-30M	100GbE QSFP-to-QSFP active optical cables, 30 m

Software Licenses (non-SONiC versions only)

QFX5K-C2-PFL	QFX5K Class 2 Premium Feature License
QFX5K-C2-AFL	QFX5K Class 2 Advanced Feature License

Model Number	Description
S-QFX5K-C3-A1-3/5 (Flex licensing)	3/5 Year Advanced 1 License for QFX5100 and QFX5200 Class 3 Products
S-QFX5K-C3-A1-P (Flex licensing)	Perpetual Advanced 1 License for QFX5100 and QFX5200 Class 3 products
S-QFX5K-C3-A2-3/5 (Flex licensing)	3/5 Year Advanced 2 License for QFX5100 and QFX5200 Class 3 Products
S-QFX5K-C3-A2-P (Flex licensing)	Perpetual Advanced 2 License for QFX5100 and QFX5200 Class 3 products
S-QFX5K-C3-P1-3/5 (Flex licensing)	3/5 Year Premium 1 License for QFX5100 and QFX5200 Class 3 Products
S-QFX5K-C3-P1-P (Flex licensing)	Perpetual Premium 1 License for QFX5100 and QFX5200 Class 3 products

Software Licenses (SONiC versions only)

Juniper cRPD is available for QFX5210-64C-S SONiC systems. For pricing information, please contact your Juniper sales representative.

About Juniper Networks

Juniper Networks brings simplicity to networking with products, solutions, and services that connect the world. Through engineering innovation, we remove the constraints and complexities of networking in the cloud era to solve the toughest challenges our customers and partners face daily. At Juniper Networks, we believe that the network is a resource for sharing knowledge and human advancement that changes the world. We are committed to imagining groundbreaking ways to deliver automated, scalable, and secure networks to move at the speed of business.

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: +31.0.207.125.700

