

QFX5200 SWITCH

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. Additionally, the QFX5200-32C-S platform includes support for SONiC, continuing Juniper's commitment to disaggregation and 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® 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 multitier, 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, 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:** 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 4.8 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:** Offering 48 ports of 25GbE and delivering 3.6 Tbps of line-rate switching, the QFX5200-48Y 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.

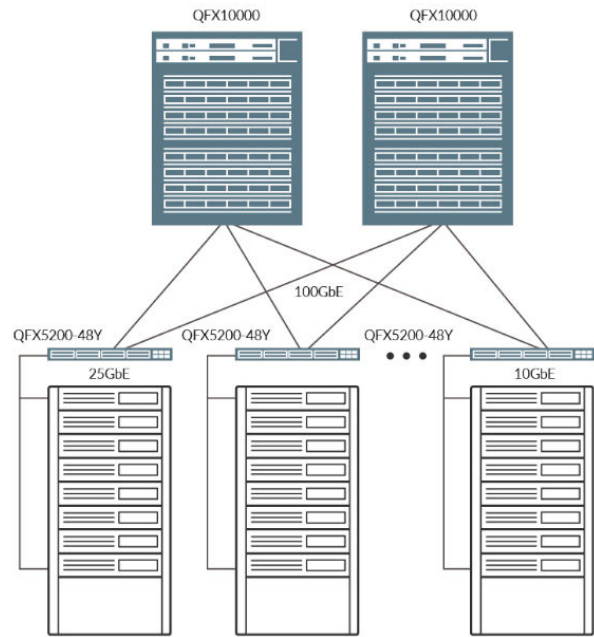


Figure 1: QFX5200-48Y and QFX10000 leaf-spine deployment

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.

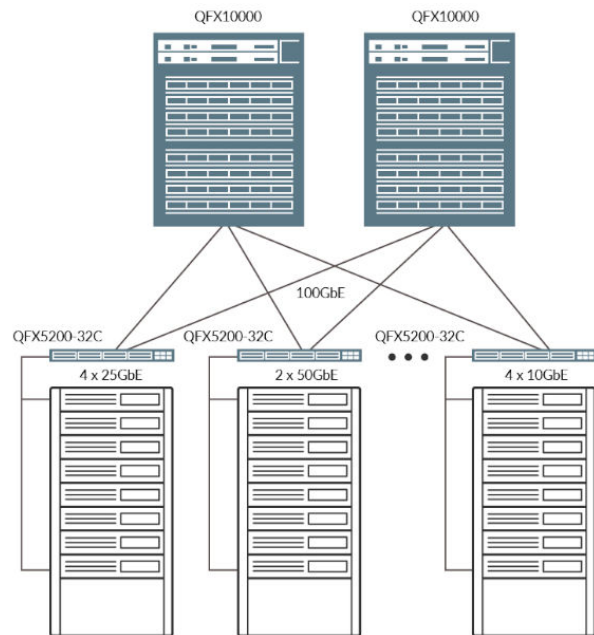


Figure 2: QFX5200-32C and QFX10000 leaf-spine deployment with centralized gateway

QFX5200 Deployment Options

The following table depicts some of the many QFX5200 deployment options, including top-of-rack access as well as spine-and-leaf configurations in a centralized or distributed gateway architecture.

Port Combination	Switch	Deployment
48x10GbE + 6x40GbE	QFX5200-48Y	10GbE access
48x25GbE + 6x100GbE	QFX5200-48Y	25GbE access
32x100GbE	QFX5200-32C	Spine
128x25GbE	QFX5200-32C	Leaf
64x25GbE + 16x100GbE	QFX5200-32C	25GbE access
96x10GbE + 8x100GbE	QFX5200-32C	10GbE access

Figure 1 shows QFX5200-48Y 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.

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.

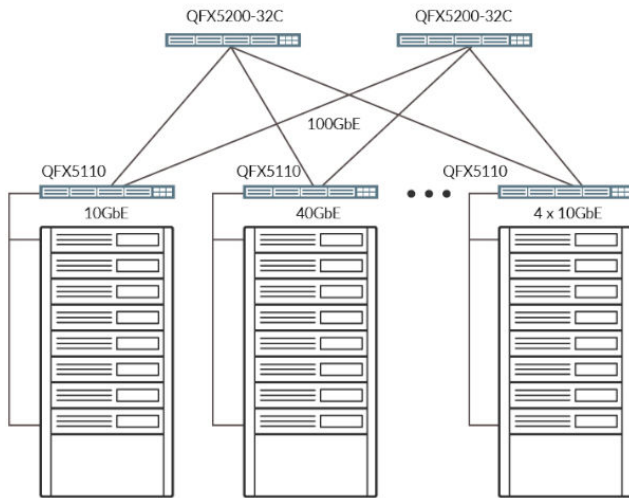


Figure 3: QFX5110 and QFX5200-32C leaf-spine deployment with distributed gateway

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 leaves 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, Juniper 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 Juniper routers that power the world's largest and most complex networks.

By using a common operating system, Juniper 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

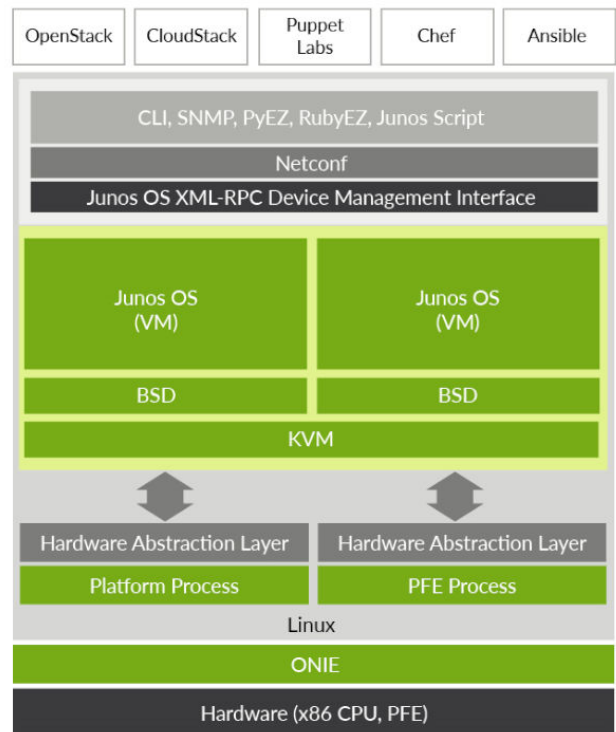


Figure 4: Junos OS software infrastructure

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

Monitoring and Analytics with Juniper Contrail Insights

Juniper Networks Contrail® Insights, a network and device monitoring platform, provides comprehensive visualization and smart analytics for the QFX5200 switches. Powered by machine learning and artificial intelligence, Contrail Insights automatically learns network and device performance profiles, detects faults, and facilitates preventative remediation in real time at scale.

Based on a distributed policy and analytics architecture, Contrail Insights seamlessly collects and analyzes telemetry in real-time from the QFX5200 platforms via JTI and OpenConfig, generating events and alert notifications and providing actionable insights to the operator.

Contrail Insights delivers the following features and benefits:

- **Real-time metrics and alarms:** Contrail Insights collects and analyzes comprehensive metrics in real time for monitored networks and devices, displaying them in the Contrail Insights dashboard as real-time charts and graphs. With Contrail Insights, an operator can also configure alarms for individual or groups of metrics that are triggered when specific values or conditions are met, or when anomalies occur, notifying the operator in real time. Contrail Insights software also plots these alarms on both real-time and historical charts, giving the operator a holistic view of current status as well as prevailing trends.
- **SLA Monitor:** The Contrail Insights dashboard includes an SLA Monitor pane that provides operators with an overview indicating the current state of the network and devices. The SLA status indicates whether the network and the device are currently operating outside a user-defined performance threshold, while risk predicts whether the device may be unhealthy in the future based on historical trends.

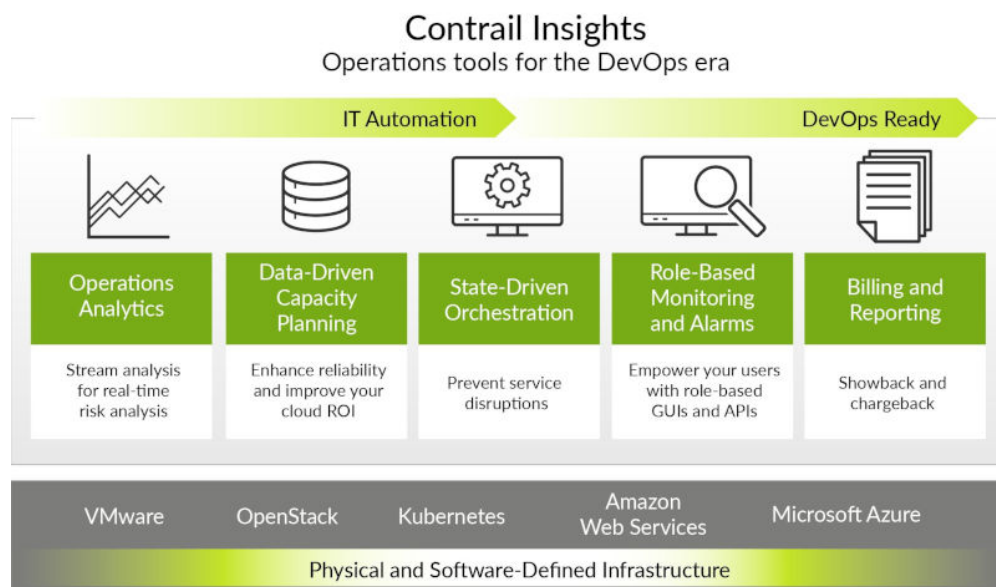


Figure 5: Contrail Insights overview

- **Notifications:** Contrail Insights can be configured to notify external systems when events and alarms are triggered. For each event or alarm, Contrail Insights can post a structured description of the event as a JSON payload to an external HTTP endpoint. These notifications can initiate an action or workflow, whether corrective, preventive, or otherwise, to keep the network and devices operating in an optimal state. For operator convenience, Contrail Insights comes preconfigured with PagerDuty, ServiceNow, and Slack, some of the industry's most popular incident management and collaboration platforms.
- **Reports:** Contrail Insights reports allow users to analyze how networks and devices are being consumed over time. The dashboard displays reports in both graphical and tabular formats. Users may also download the data as an HTML-formatted report, raw comma-separated value (CSV) file, or JSON-formatted statistics for further analysis.

Contrail Networking

The QFX5200 switches can be managed with Juniper Networks Contrail Networking, a fabric management and SDN overlay solution that provides full lifecycle management for both data center and data center interconnect infrastructures. Contrail Networking automates data center operations so users can focus on revenue-generating services rather than operational networking procedures.

Contrail Networking works with any standard IP Clos architecture, performing configuration operations across virtual networks and physical devices by leveraging open protocols such as NETCONF/

RPC. Using a centralized controller with a management interface, Contrail Networking provides customizable configuration templates that adapt to a vast array of architectures and use cases to ease network and device operations. Contrail Networking abstracts network connectivity using virtual overlays connecting cloud-native workloads with traditional workloads running on bare-metal servers attached to the QFX5200.

Contrail Networking enables the provisioning and automation of data center fabrics and data center interconnect using a common, consistent data model for configuring devices. Key capabilities include:

- Infrastructure, multicloud interconnect, and service lifecycle management
- Open, scalable, and standards-based protocol support: EVPN/VXLAN, BGP, NETCONF/RPC, sFlow/gRPC/SNMP
- Autodiscovery, Zero-Touch Provisioning (ZTP), Zero-Touch Replacement (ZTR), and upgrades of data center fabrics including QFX5200 switches
- Overlay and underlay management from a single orchestration layer
- Seamless connection between legacy and cloud-native workloads
- Network and device topology discovery with role configuration
- BGP EVPN across QFX Series switches and virtual routers

Contrail Networking and Contrail Insights are fundamental building blocks of Contrail Enterprise Multicloud, a data center software suite to connect, orchestrate, monitor, and secure multiclouds.

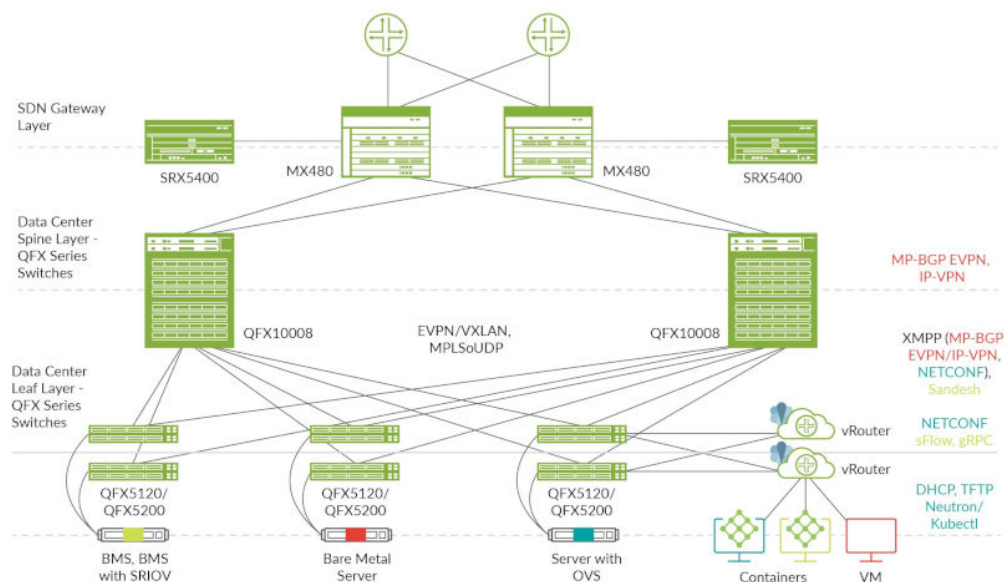


Figure 6: Contrail Networking as a fabric management and SDN overlay for data center infrastructure.

Junos Space Network Director

The QFX5200 can be managed through Juniper Networks Junos Space® Network Director*, a next-generation network management solution that lets you visualize, analyze, and control the entire enterprise network—data center and campus, physical and virtual—through a single pane of glass. Network Director incorporates sophisticated analytics for real-time intelligence, trended monitoring, and automation to increase agility and ensure faster rollout and activation of services.

For cloud deployments, Network Director provides a set of REST APIs that enable on-demand and dynamic network services by simplifying the consumption of services for multitenant environments. With third-party cloud orchestration tool integration, the Network Director API enables automation and provisioning of L2, L3, and security services in the data center without the need for manual operator intervention.

Features and Benefits

- **Automation and Programmability:** The QFX5200 supports numerous network automation features, including operations and event scripts, ZTP, and Juniper 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-S Hardware with SONiC Support

The QFX5200-32C-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 QFX5200-32C-S, which plugs seamlessly into a unified SONiC network infrastructure, leverages the existing QFX5200-32C 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 Layer 2, Layer 3, MPLS, EVPN, and VXLAN, among others. 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 Daemon (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, the cRPD enables high-performance telemetry, automation, and programmability in a lightweight deployment.



QFX5200-32C



QFX5200-48Y

Specifications

Hardware

Table 1. QFX5200 System Capacity

	QFX5200-32C	QFX5200-48Y
System throughput	Up to 6.4 Tbps (bi-directional)	3.6 Tbps (bi-directional)
Forwarding capacity	Up to 2.4 Bpps	2.1 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
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
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	<ul style="list-style-type: none"> 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 	<ul style="list-style-type: none"> 1 RJ-45 10/100/1000 BASE-T for management 1 RS-232 console port 1 USB 2.0 port
Power	<ul style="list-style-type: none"> Redundant (1+1) hot-pluggable 850 W AC/DC power supplies 110-240 V single phase AC power -36 to -72 V DC power 	<ul style="list-style-type: none"> Redundant (1+1) hot-pluggable 650 W AC/DC power supplies
Cooling	<ul style="list-style-type: none"> Front-to-back and back-to-front cooling Redundant (N+1) hot-pluggable fan modules with variable speed to minimize power draw 	<ul style="list-style-type: none"> Front-to-back and back-to-front cooling Redundant (N+1) hot-pluggable fan modules
Total packet buffer	16 MB	22 MB
Warranty	Juniper standard one-year warranty	Juniper standard one-year warranty

Software (Non SONiC Versions)

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

- 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)

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, PIM-Bidir*
- 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: Policier, policier 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)*

Server Virtualization Management and SDN-Related Protocols

- IEEE 802.1Qbg (VEPA hairpin forwarding)

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

- Contrail Networking
- Junos Space Network Director*
- 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
- 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

Supported MIBs

- RFC 155 SMI
- RFC 1157 SNMPv1
- RFC 1212, RFC 1213, RFC 1215 MIB-II, Ethernet-Like MIB and TRAPs

Environmental Ranges

Parameters	QFX5200-32C	QFX5200-48Y-AFO	QFX5200-48Y-AFI
Operating temperature	32° to 104° F (0° to 40° C)	32° to 104° 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

Maximum Thermal Output

Parameters	QFX5200-32C	QFX5200-48Y
Maximum power draw	480 W (AC, DC), 1638 BTU/hr	430 W (AC, DC), 1467 BTU/hr
Typical power draw	380 W (AC, DC), 1296 BTU/hr	382 W (AC, DC), 1303 BTU/hr

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, Authorisation and Restriction of Chemicals (REACH)



Waste Electronics and Electrical Equipment (WEEE)



Recycled material



80 Plus Silver PSU Efficiency

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. Support and services for the SONiC-enabled QFX5200-32C-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

Product	Description	Product	Description
Switch Hardware			
QFX5200-32C-AFI	QFX5200 (hardware only; software services sold separately), 32 QSFP+/QSFP28 ports, redundant fans, 2 AC power supplies, back-to-front airflow	QFX-SFP-10GE-SR	SFP+ 10GBASE-SR 10GbE optics, 850 nm for up to 300 m transmission on multimode fiber-optic (MMF) (for management port)
QFX5200-32C-AFO	QFX5200 (hardware only; software services sold separately), 32 QSFP+/QSFP28 ports, redundant fans, 2 AC power supplies, front-to-back airflow	QFX-SFP-10GE-LR	SFP+ 10GBASE-LR 10GbE optics, 1310 nm for 10 km transmission on single-mode fiber-optic (SMF) (for management port)
QFX5200-32C-DC-AFI	QFX5200 (hardware only; software services sold separately), 32 QSFP+/QSFP28 ports, redundant fans, 2 DC power supplies, back-to-front airflow	QFX-SFP-DAC-1M	SFP+10 GbE direct attach copper (DAC) (twinax copper cable) 1 m for QFX5200-48Y
QFX5200-32C-DC-AFO	QFX5200 (hardware only; software services sold separately), 32 QSFP+/QSFP28 ports, redundant fans, 2 DC power supplies, front-to-back airflow	QFX-SFP-DAC-3M	SFP+10 GbE DAC (twinax copper cable) 3 m for QFX5200-48Y
QFX5200-32C-DCSAFI	QFX5200, 32 QSFP+ ports, redundant fans, 2 DC power supplies, back-to-front airflow, with SONiC	QFX-SFP-DAC-5M	SFP+10 GbE DAC (twinax copper cable) 5 m for QFX5200-48Y
QFX5200-32C-DCSAFO	QFX5200, 32 QSFP+ ports, redundant fans, 2 DC power supplies, front-to-back airflow, with SONiC	JNP-SFP-25G-SR	25GBASE-SR SFP28 module for MMF for QFX5200-48Y
QFX5200-32C-SAFI	QFX5200, 32 QSFP+ ports, redundant fans, 2 AC power supplies, back-to-front airflow, with SONiC	JNP-SFP-25G-DAC-1M	SFP 25GBASE DAC cable 1 m for QFX5200-48Y
QFX5200-32C-SAFO	QFX5200, 32 QSFP+ ports, redundant fans, 2 AC power supplies, front-to-back airflow, with SONiC	JNP-SFP-25G-DAC-3M	SFP 25GBASE DAC cable 3 m for QFX5200-48Y
QFX5200-32C-S-CHAS	QFX5200-32C chassis with SONiC as default NOS.	JNP-SFP-25G-DAC-5M	SFP 25GBASE DAC cable 5 m for QFX5200-48Y
JPSU-850W-AC-AFI	Juniper 850 W AC power supply for QFX5100-96S and QFX5200-32C (power supply-side airflow intake)	JNP-QSFP-40GE-IR4	QSFP+ 40GBASE-IR4 40-gigabit optics, 1310 nm for up to 2 km transmission on SMF
JPSU-850W-AC-AFO	Juniper 850 W AC power supply for QFX5100-96S and QFX5200-32C (power supply-side airflow exhaust)	JNP-QSFP-40G-LR4	QSFP+ 40GBASE-LR4 40-gigabit optics, 1310 nm for up to 10 km transmission on SMF
JPSU-850W-DC-AFI	DC 850 W power supply, back-to-front airflow for QFX5100-96S and QFX5200-32C	JNP-QSFP-40G-LX4	QSFP+ 40GBASE-LX4 40-gigabit optics, 100 m (150 m) with OM3 (OM4) duplex MMF fiber
JPSU-850W-DC-AFO	DC 850 W power supply, front-to-back airflow for QFX5100-96S and QFX5200-32C	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)
QFX5200-32C-FANAFI	QFX5200-32C fan module, back-to-front airflow	JNP-QSFP-4X10GE-LR	QSFP+ 40GBASE optics, up to 10 km transmission on parallel single mode (4X10GbE long reach up to 10 km)
QFX5200-32C-FANAFO	QFX5200-32C fan module, front-to-back airflow	QFX-QSFP-40G-ESR4	QSFP+ 40GBASE-ESR4 40-gigabit optics, 300 m (400 m) with OM3 (OM4) MMF
QFX5200-48Y-AFI	QFX5200-48Y (base software services included), 48 SFP+/SFP28 ports, redundant fans, 2 AC power supplies, back-to-front airflow	QFX-QSFP-40G-SR4	QSFP+ 40GBASE-SR4 40-gigabit optics, 850 nm for up to 150 m transmission on MMF
QFX5200-48Y-AFO	QFX5200-48Y (base software services included), 48 SFP+/SFP28 ports, redundant fans, 2 AC power supplies, front-to-back airflow	JNP-40G-AOC-20M	40GbE active optical cable for 20 m
QFX5200-48Y-DC-AFI	QFX5200-48Y (base software services included), 48 SFP+/SFP28 ports, redundant fans, 2 DC power supplies, back-to-front airflow	JNP-40G-AOC-30M	40GbE active optical cable for 30 m
QFX5200-48Y-DC-AFO	QFX5200-48Y (base software services included), 48 SFP+/SFP28 ports, redundant fans, 2 DC power supplies, front-to-back airflow	JNP-QSFP-DAC-10MA	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 10 m active
QFX520048Y-APSU-AO	Juniper 650 W AC power supply for QFX5200-48Y (power supply-side airflow intake)	JNP-QSFP-DAC-7MA	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 7 m active
QFX520048Y-APSU-AI	Juniper 650 W AC power supply for QFX5200-48Y (power supply-side airflow exhaust)	JNP-QSFP-DAC-5M	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 5 m passive
QFX520048Y-DPSU-AO	DC 650 W power supply, front-to-back airflow for QFX5200-48Y	QFX-QSFP-DAC-3M	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 3 m passive
QFX520048Y-DPSU-AI	DC 650 W power supply, back-to-front airflow for QFX5200-48Y	QFX-QSFP-DAC-1M	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 1 m passive
QFX520048Y-FAN-AO	QFX5200-48Y fan module, front-to-back airflow	EX-QSFP-40GE-DAC-50CM	40GbE QSFP+ 0.5 m direct-attach
QFX520048Y-FAN-AI	QFX5200-32C fan module, back-to-front airflow	JNP-QSFP-DACBO-10M	QSFP+ to SFP+ 10GbE direct attach breakout copper (twinax copper cable) 10 m active
EX-4PST-RMK	Adjustable 4-post rack-mount kit for EX Series, QFX5100, and QFX5200	JNP-QSFP-DACBO-7MA	QSFP+ to SFP+ 10GbE direct attach breakout copper (twinax copper cable) 7 m active
QFX520048Y-RKMT	4 post rack-mount for QFX5200-48Y	JNP-QSFP-DACBO-5MA	QSFP+ to SFP+ 10GbE direct attach breakout copper (twinax copper cable) 5 m active
Optics and Transceivers		QFX-QSFP-DACBO-3M	QSFP+ to SFP+ 10GbE direct attach breakout copper (twinax copper cable) 3 m
QFX-SFP-10GE-USR	SFP+ 10GbE Ultra Short Reach optics, 850 nm for 10 m on OM1, 20 m on OM2, 100 m on OM3 multimode fiber (for management port)	QFX-QSFP-DACBO-1M	QSFP+ to SFP+ 10GbE direct attach breakout copper (twinax copper cable) 1 m
		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-QSFP-100G-PSM4	100GbE PSM4 optics
		JNP-QSFP-100G-CWDM	100GbE CWDM4 optics

Product	Description
JNP-100G-AOC-1M	100GbE active optical cable for 1 m
JNP-100G-AOC-15M	100GbE active optical cable for 15 m
JNP-100G-AOC-20M	100GbE active optical cable for 20 m
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 (SONiC Version)	
QFX-QSFP-40G-SR4	QSFP+ 40GBASE-SR4 40 Gigabit Optics, 850nm for up to 150m transmission on MMF
JNP-QSFP-40G-LR4	QSFP+ 40GBASE-LR4 40 Gigabit Optics for up to 10km transmission on SMF
JNP-QSFP-100G-SR4	QSFP28 100GBASE-SR4 Optics for up to 100m transmission over parallel MMF
JNP-QSFP-100G-LR4	QSFP28 100GBASE-L4 Optics for up to 10km transmission over serial SMF
JNP-QSFP-4x10GE-LR	QSFP+ 40GBASE Optics, up to 10km Transmission on parallel single mode (4X10GE LR up to 10km)
JNP-100G-AOC-1M	100GbE active optical cable for 1 m
JNP-100G-AOC-3M	100GbE active optical cable for 3 m
JNP-100G-AOC-5M	100GbE active optical cable for 5 m
JNP-100G-DAC-1M	100GbE DAC Twinex Copper Cable for 1 m
JNP-100G-DAC-3M	100GbE DAC Twinex Copper Cable 3 m
JNP-100G-DAC-5M	100GbE DAC Twinex Copper Cable 5 m
QFX-QSFP-DAC-1M	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 1 m passive
JNP-100G-4x25G-1M	100G QFP28 to SFP28 4x25G Direct Attach Copper Breakout Twinex Copper Cable 1 m
JNP-100G-4x25G-3M	100G QFP28 to SFP28 4x25G Direct Attach Copper Breakout Twinex Copper Cable 3 m
Software Licenses	
QFX5000-35-JBS	Base Services license for QFX5200-32C
QFX5000-35-JAS	Advanced Services license for QFX5200-32C
QFX5000-35-JPS	Premium services license for QFX5200-32C
QFX5K-C1-PFL	QFX5000 Class 1 Premium Feature License for QFX5200-48Y
QFX5K-C1-AFL	QFX5000 Class 1 Advanced Feature License for QFX5200-48Y

Software Licenses (SONiC Versions Only)

Juniper cRPD is available for QFX5200-32C-S SONiC systems. For pricing information, please contact your Juniper sales representative.

* Not supported on QFX5200-48Y at FRS

Note: Base Services software license is not included with the QFX5200-32C hardware and is sold separately.

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

