![JUNIPER NETWORKS]

QFX5100 ETHERNET SWITCH

Product Overview

The Juniper Networks QFX5100 line of access and aggregation switches delivers low latency, flexible deployment options, and rich Layer 2 and Layer 3 features, making it the industry's most nimble family of switches.

Featuring Topology-Independent In-Service Software Upgrades (TISSU) and Insight Technology for microburst monitoring and hotspot statistics, the QFX5100 is designed for the most demanding, high-performance data center deployments. Highly flexible, the QFX5100 can be deployed in mixed 1GbE, 10GbE, and 40GbE environments within Juniper Virtual Chassis, Virtual Chassis Fabric and QFabric architectures, as well as within industry-standard architectures such as spine and leaf and Layer 3 fabrics.

Product Description

The highly flexible, high-performance Juniper Networks® QFX5100 line of switches provides the foundation for today's and tomorrow's dynamic data center. As a critical enabler for IT transformation, the data center network supports cloud and software-defined networking (SDN) adoption, as well as rapid deployment and delivery of applications. Mission-critical applications, network virtualization, and integrated or scale-out storage are driving the need for more adaptable networks. With its diverse set of deployment options, including fabric, Layer 3, and spine and leaf, the QFX5100 is the universal building block for data center switching architectures, enabling users to easily adapt as requirements change over time.

The QFX5100 line includes 10GbE and 40GbE fixed-configuration options with rich Layer 2, Layer 3, and MPLS features. The QFX5100 runs the same reliable, high-performance Juniper Networks Junos® operating system that is used by Juniper Networks QFabric™ family of products, EX Series Ethernet Switches, Juniper routers, and Juniper Networks SRX Series Services Gateways, ensuring a consistent implementation and operation of control plane features across the entire Juniper infrastructure.

Featuring a built-in Insight Technology capability, the QFX5100 provides valuable performance and troubleshooting data via microburst monitoring and hotspot statistics. The QFX5100 can also be used as a high-performance QFabric Node in Juniper's high-scale QFabric System, deployed in Juniper's new Virtual Chassis Fabric switching architecture, and installed in Juniper's existing Virtual Chassis switching architecture.

Architecture and Key Components

The QFX5100 includes the following capabilities. Please refer to the Specifications section for features that are currently shipping vs. those that will be available in a future software release.

QFX5100 Highlights

- High-density configurations—72 ports of 10GbE in a 1 U platform; 96 ports of 10GbE in a 2 U platform; and 32 ports of 40GbE in a 1 U platform
- Up to 2.56 Tbps Layer 2 and Layer 3 performance and as low as 550 ns latency*
- 1.5 GHz dual-core Intel CPU with 8 GB memory and 32 GB SSD storage
- High availability with the industry's only Topology-Independent In-Service Software Upgrade (TISSU)
- Rich automation capabilities with support for Python, Chef, Puppet, and zero-touch provisioning (ZTP)

* 40GbE-to-40GbE port on QFX5100-24Q only.

- Designed-in Insight Technology software capability for microburst monitoring and reporting, as well as hotspot statistics and visibility
- Plug-and-play deployable architectures such as Virtual Chassis, Virtual Chassis Fabric, and QFabric System with simplified management and operation
- Support for standards-based network virtualization protocols such as Virtual Extensible LAN (VXLAN), Open vSwitch Database (OVSDB) protocol, Network Virtualization using Generic Routing Encapsulation (NVGRE), and integration with Juniper Networks Contrail and VMware NSX SDN controllers
- Advanced Junos OS features such as BGP Add-path, MPLS, L3 VPN, and IPv6 6PE
- Investment protection enabled by compatibility with Juniper Networks EX4300 Ethernet Switch, QFX3500 Switch, and QFX3600 Switch in integrated fabric solutions

QFX5100 Switch Models

The QFX5100 line includes two compact 1 U models and one 2 U model, each providing 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 control plane running on all QFX5100 models is further enhanced with a powerful 1.5 GHz dual-core Intel CPU with 8 GB of memory and 32 GB SSD storage.

QFX5100-48S: Compact 1 U 10GbE data center access switch with 48 small form-factor pluggable and pluggable plus (SFP/SFP+) transceiver ports and six quad SFP+ (QSFP+) ports with an aggregate throughput of 1.44 Tbps or 1.08 Bpps per switch.

QFX5100-24Q: Compact 1 U high-density 40GbE data center access and aggregation switch starting at a base density of 24 QSFP+ ports with the option to scale to 32 QSFP+ ports with two four-port expansion modules. All 32 ports support wire-speed performance with an aggregate throughput of 2.56 Tbps or 1.44 Bpps per switch.

QFX5100-96S: Compact 2 U high-density 10GbE aggregation switch with 96 SFP+/SFP and eight QSFP+ ports with an aggregate throughput of 2.56 Tbps or 1.44 Bpps per switch.

For added flexibility, the 40GbE ports on all QFX5100 switches can be used as 4x10GbE ports using QSFP+ to SFP+ direct attach copper (DAC) or QSFP+ to SFP+ fiber splitter cables and optics. With splitter cables, the QFX5100-48S supports up to 72 10GbE ports, while the QFX5100-24Q and QFX5100-96S support up to 104 10GbE ports.

Junos OS

All three high-performance QFX5100 switches run Junos OS, Juniper's powerful and robust network operating system. Key Junos OS features that enhance the functionality and capabilities of the QXF5100 include:

- Software modularity with process modules running independently in their own protected memory space and with the ability to do process restarts
- Uninterrupted routing and forwarding with features such as nonstop routing (NSR) and nonstop bridging (NSB)
- Commit and rollback functionality that ensures error-free network configurations
- A powerful set of scripts for on-box problem detection, reporting, and resolution

Junos Space Network Director

The QFX5100 switches can be managed through Junos Space Network Director, a next-generation network management solution that allows users to visualize, analyze, and control the entire enterprise network—data center and campus, physical and virtual, wired and wireless—through a single pane of glass. Network Director incorporates sophisticated analytics for real-time intelligence, trended monitoring, and automation to increase agility as well as 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 Layer 2, Layer 3, and security services in the data center without the need for manual operator intervention.

Virtual Chassis Technology

The QFX5100 supports Juniper Networks' unique Virtual Chassis technology, which enables up to 10 interconnected switches to be managed and operated as a single, logical device with a single IP address. Virtual Chassis technology enables enterprises to separate physical topology from logical groupings of endpoints and, as a result, provides efficient resource utilization. All QFX5100 switches can participate in the same Virtual Chassis configuration, along with any combination of EX4300, QFX3500, and QFX3600 Ethernet switches, delivering highly flexible and scalable configuration options for data center deployments.

Virtual Chassis Fabric Switching Architecture

Existing Virtual Chassis technology is further scaled and enhanced to support a spine-and-leaf topology that is ideal for high-performance and low-latency data center deployments. In its first instance, this topology, called Virtual Chassis Fabric, enables up to 20 QFX5100 switches to be deployed in a spine-and-leaf configuration, with two to four QFX5100s in the spine and up to 18 QFX5100s as leaf nodes. This architecture provides any-rack-to-any-rack deterministic throughput and less than 2 microseconds of latency, while significantly simplifying network operations through a single point of management. EX4300, QFX3500, and QFX3600 switches can also act as leaf nodes in a Virtual Chassis Fabric deployment, providing support for data centers with mixed 1GbE, 10GbE, and 40GbE servers as well as environments transitioning from 1GbE to 10GbE or 40GbE servers while protecting existing EX4300, QFX3500, and QFX3600 investments.

QFabric System

The QFX5100-48S can also act as a QFabric Node device in a QFX3000-M and QFX3000-G QFabric System. The QFX5100-48S is fully interoperable with existing QFX3500 and QFX3600 QFabric Nodes, providing backward compatibility and investment protection. With a mixture of QFX3500, QFX3600, and QFX5100 Nodes, a QFabric System can scale on a node by node basis without being limited by the lowest common denominator.

Data Center Deployments

Today's data centers are built with high-performance, small form-factor, multicore blade and rack servers. The greater compute capacity and server densities enabled by these devices is increasing traffic levels, creating the need for high-speed, low-latency, storage- and I/O-converged networking solutions that can maximize performance for physical servers, virtual servers, and storage in the data center.

The QFX5100 line offers a portfolio of switches that deliver the low-latency, lossless, high-density 10GbE and 40GbE interfaces with the Fibre Channel over Ethernet (FCoE) transit switch functionality demanded by today's data center. All QFX5100 models are designed to consume the lowest power possible while optimizing space, thereby reducing data center operating costs. Flexible airflow direction options enable the QFX5100 to support back-to-front and front-to-back cooling, ensuring consistency with server designs for hot- and cold-aisle deployments.

Data Center Server Access

The QFX5100 line of switches is ideal for top-of-rack deployments of various densities and speeds.

The compact 1 U QFX5100-48S offers 48 ports of native 10GbE for server connectivity and up to six 40GbE ports for uplink connectivity, providing very low oversubscription of 2:1 from access to aggregation. Each 40GbE port can be further broken out into four 10GbE ports, providing additional 10GbE for server connectivity.

The 2 U QFX5100-96S is a high-density native 10GbE data center access switch that is ideal for deployments with a high degree of application redundancy. The highly efficient QFX5100-96S offers 96 ports of native 10GbE for server connectivity and eight 40GbE ports for uplink connectivity, providing very low oversubscription of 3:1 from access to aggregation.

The QFX5100-24Q offers 24 QSFP+ ports in the base unit with an option to scale up to 32 QSFP+ ports with two four-port expansion modules. Each QSFP+ port can be used for native 40GbE server port connectivity or channelized into four 10GbE ports for 10GbE server connectivity, offering maximum flexibility and investment protection for data center growth.

Each of the QFX5100 switches can operate in both cut-through and store-and-forward modes, delivering sustained wire-speed switching with sub-microsecond latency and low jitter for any packet size (including jumbo frames) in both modes. All QFX Series switches support extensive Layer 2 features, enabling the device to support high-density 10GbE/FCoE Layer 2 access deployments. With features such as multi-chassis link aggregation group (MC-LAG), the QFX5100 supports active/active server dual homing and can utilize full bisectional bandwidth from server to switch. When the QFX5100 is deployed in the aggregation layer, MC-LAG at the access and aggregation layers provides maximum resiliency and full Layer 2 multi-pathing in the network (see Figures 1 and 2).

The Junos operating system features the most advanced and robust routing capabilities in the industry. All QFX5100 switches include support for RIP and OSPF for both IPv4 and IPv6 in the base software. Advanced routing capabilities such as IS-IS and BGP are also supported. With additional capabilities like 64-way equal-cost multipath (ECMP) and BGP add path, QFX5100 is an ideal building block for deploying the most robust Layer 3 underlay for software-defined networks (SDN) (see Figure 3).

In addition to MC-LAG-based Layer 2 and BGP-based Layer 3 to access deployments, the QFX5100 also supports Virtual Chassis Fabric and QFabric System deployments. These fabric technologies support both Layer 2 multi-pathing and Layer 3 ECMP paths within the same network fabric while dramatically reducing the number of user configured protocols, thus eliminating the need to run protocols such as Virtual Router Redundancy Protocol (VRRP) for Layer 3 gateway load balancing and complex multicast protocols to set up replication trees. Virtual Chassis Fabric and QFabric technology dramatically reduce the complexity and cost of managing the data center network by making multiple switches appear as a single, logical device. Virtual Chassis Fabric allows up to 20 individual switches to be managed via a single CLI, while a QFabric System enables up to 128 Nodes to operate as a single device (see Figures 4 and 5).
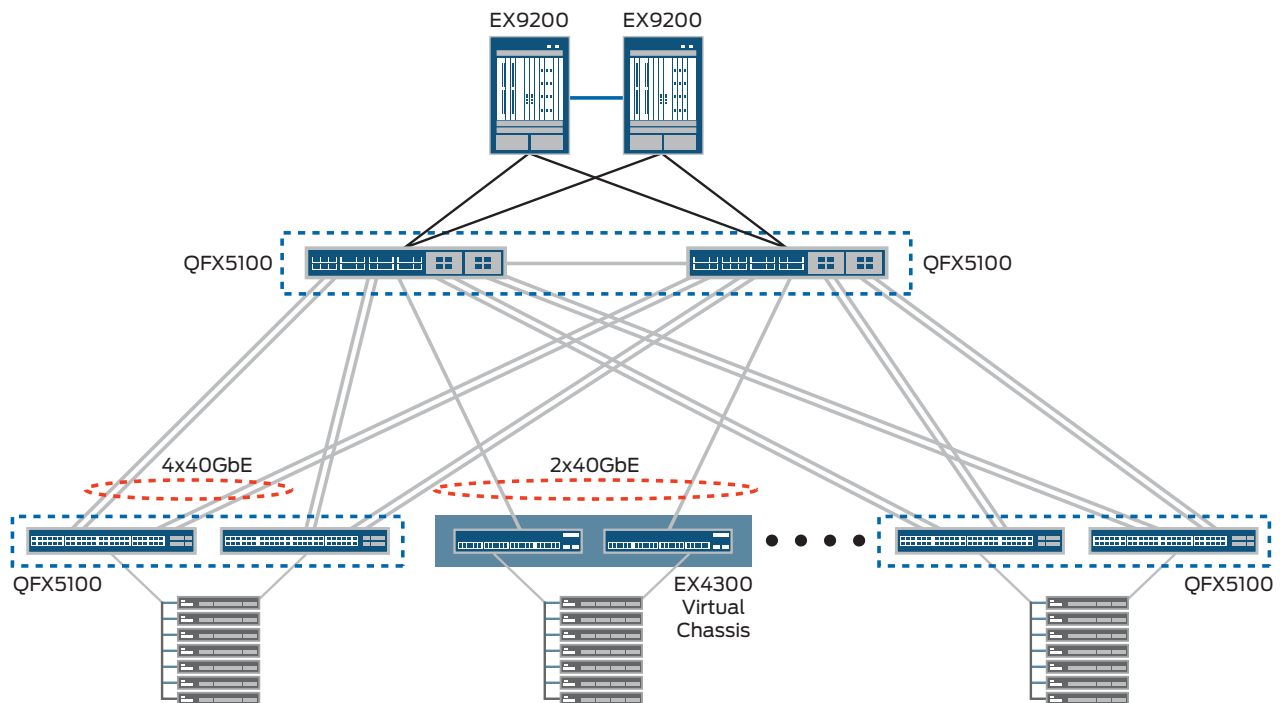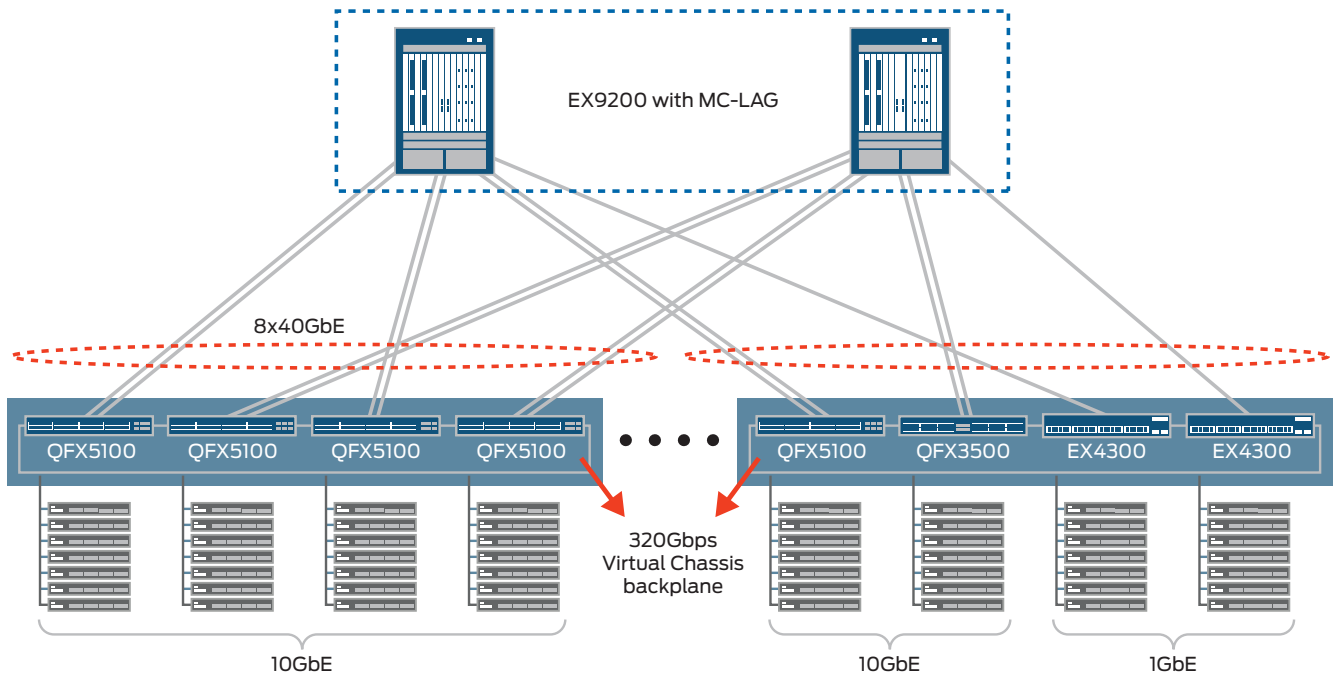

Figure 1: QFX5100 at the access layer

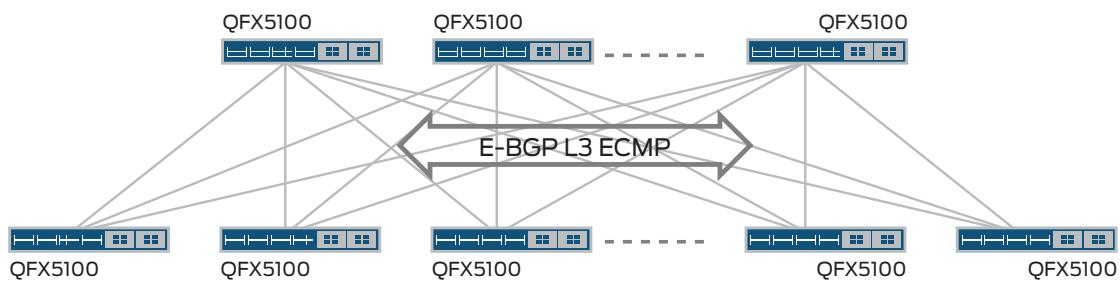Figure 2: QFX5100 at the aggregation layer

Figure 3: Larger scale data center with Layer 3 to the access

Campus Deployments

The QFX5100 offers an economical, power-efficient, and compact solution for aggregating 10GbE connections from access devices in building and campus deployments (see Figure 6). The switch's dual-speed interfaces also support environments transitioning from 1GbE to 10GbE.

The QFX5100 easily meets enterprise core switch requirements, delivering wire-speed performance on every port, full device redundancy, support for Layer 3 dynamic routing protocols such as RIP and OSPF, L2 and L3 MPLS VPNs, and a comprehensive security and quality-of-service (QoS) feature set.

Features and Benefits

- **Topology-Independent In-Service Software Upgrade (TISSU):** With its Intel core processor, the QFX5100 switches allow Junos OS to run within a virtual machine (VM) on Linux. Junos OS runs in two separate VMs in active and standby pairs; during software upgrade cycles, the switches seamlessly move to the newer software version while maintaining intact data plane traffic. This true Topology-Independent ISSU, an industry-first software upgrade feature for a fixed-configuration top-of-rack switch, is supported across all Layer 2 and Layer 3 protocols and doesn't need the support of any other switches to perform an image upgrade.

- **Automation:** The QFX5100 switches support a number of features for network automation and plug-and-play operations. Features include zero-touch provisioning, operations and event scripts, automatic rollback, and Python scripting. The switch also offers support for integration with VMware NSX Layer 2 Gateway Services, Puppet, and OpenStack.

- **Flexible Forwarding Table:** The QFX5100's Flexible Forwarding Table (FFT) allows the hardware table to be carved into configurable partitions of Layer 2 media access control (MAC), Layer 3 host, and Longest Prefix Match (LPM) tables. In a pure L2 environment, the QFX5100 supports 288,000 MAC addresses. In L3 mode, the table can support 128,000 host entries, and in LPM mode, it can support 128,000 prefixes. Junos OS provides configurable options through a command-line interface (CLI) so that each QFX5100 can be optimized for different deployment scenarios.

- **Intelligent Buffer Management:** The QFX5100 switches have a total of 12 MB shared buffers. While 25% of the total buffer space is dedicated, the rest is shared among all ports and is user
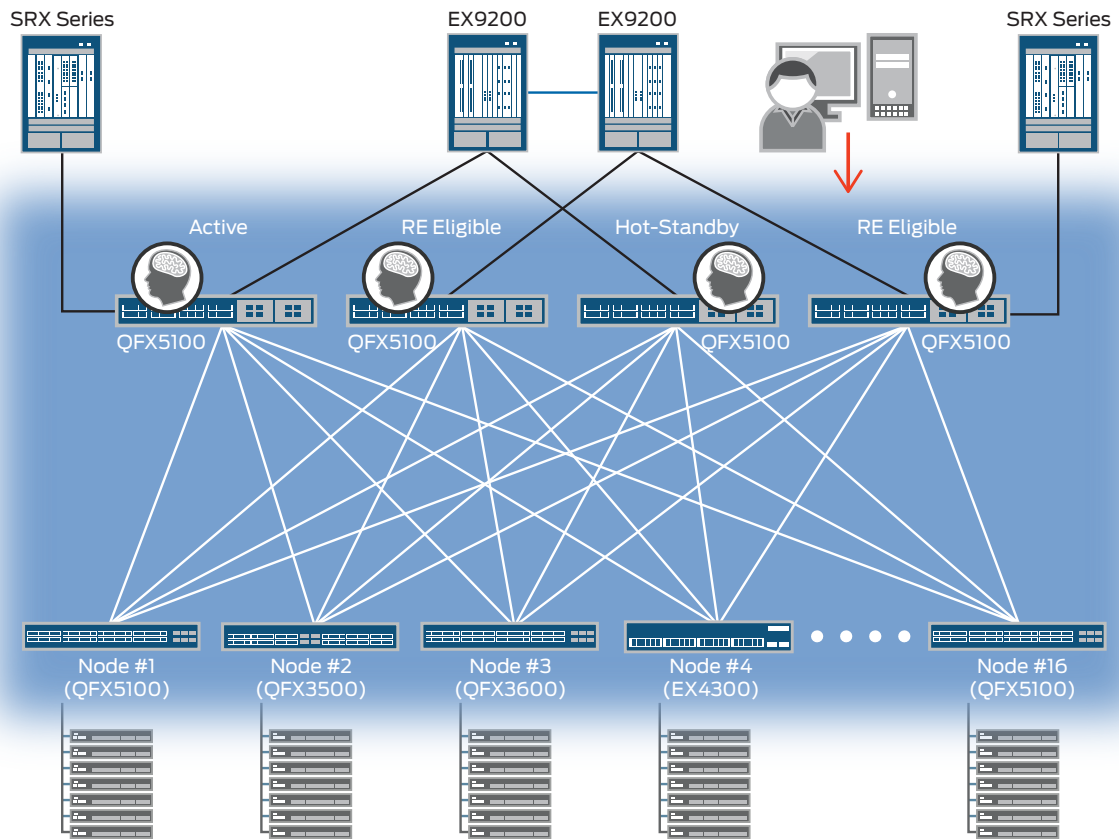
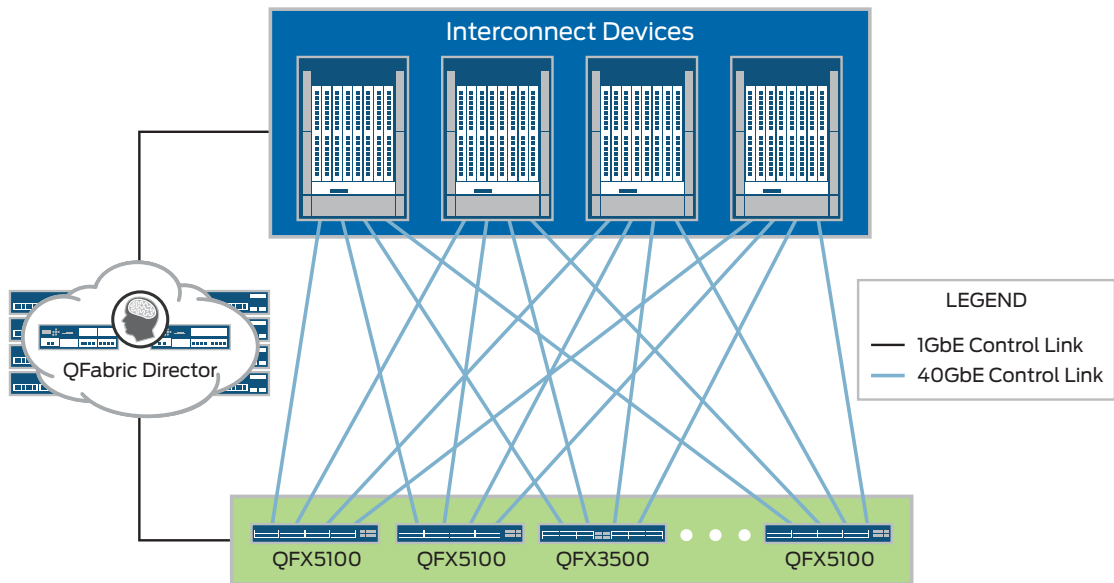Figure 4: Virtual Chassis Fabric data center deployment with a mix of 1GbE and 10GbE

Figure 5: A converged 10GbE data center network with QFabric family of products
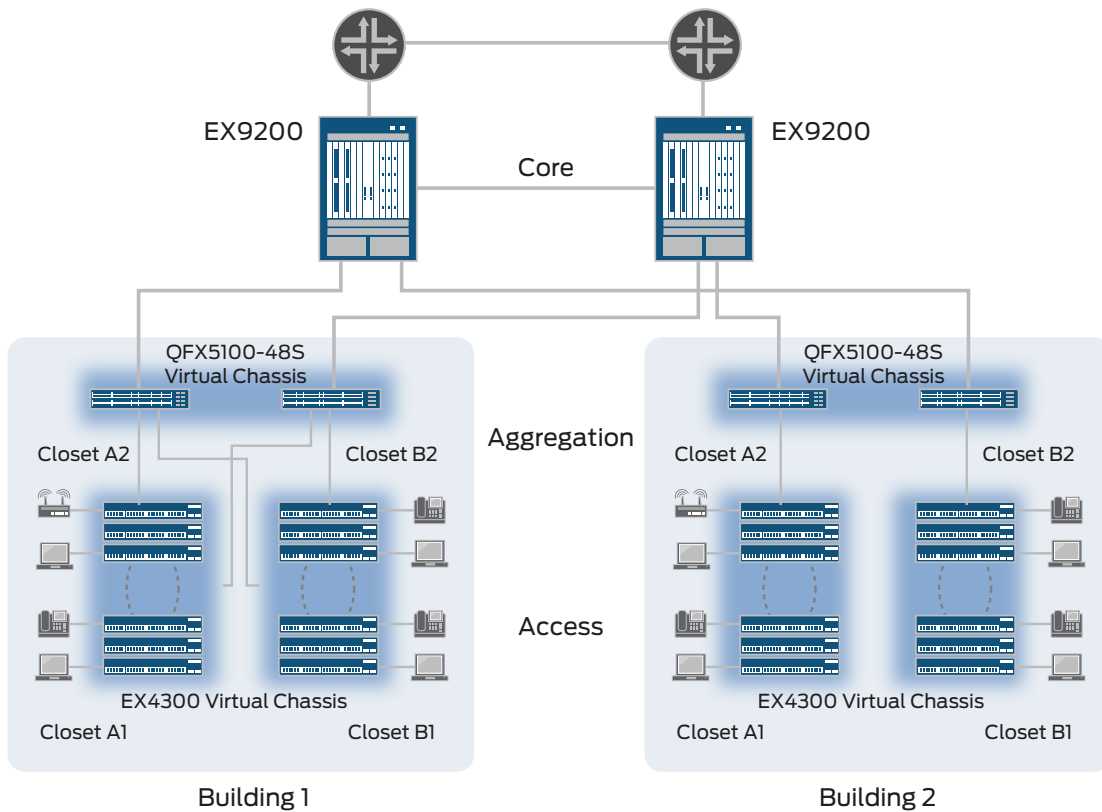
Figure 6: The QFX5100 aggregates 10GbE connections from access devices in building and campus deployments.

configurable. The intelligent buffer mechanism in the QFX5100 effectively absorbs traffic bursts while providing deterministic performance, significantly increasing performance over static allocation.

- **Insight Technology for Analytics:** The QFX5100 provides dynamic buffer utilization monitoring and reporting with an interval of 10 milliseconds to provide microburst and latency insight. It calculates both queue depth and latency, and logs messages when configured thresholds are crossed. Interface traffic statistics can be monitored at two-second granularity. The data can be viewed via CLI, system log, or streamed to external servers for more analysis. Supported reporting formats include Java Script Object Notification (JSON), CSV and TSV. These files can be consumed by orchestration systems, SDN controllers, or network management applications (such as Juniper Networks Junos Space Network Director) to make better network design decisions and identify network hotspots.

- **MPLS**: QFX5100 switches support a broad set of MPLS features, including L3 VPN, IPv6 provider edge router (6PE), RSVP traffic engineering, and LDP to allow standards-based network segmentation and virtualization. The QFX5100 can be deployed

as a low-latency MPLS label-switching router (LSR) or MPLS PE router in smaller scale environments. The QFX5100 is the industry's only compact, low-latency, high-density, low-power switch to offer an MPLS feature set.

- **FCoE**: As a Fibre Channel over Ethernet (FCoE) transit switch, the QFX5100 provides an IEEE data center bridging (DCB) converged network between FCoE-enabled servers and an FCoE-enabled Fibre Channel storage area network (SAN). The QFX5100 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 Initiation Protocol (FIP) snooping provides perimeter protection, ensuring that the presence of an Ethernet layer does not impact existing SAN security policies. FCoE link aggregation group (LAG) active/active support is available to achieve resilient (dual-rail) FCoE connectivity.

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

QFX5100-96S

QFX5100-48S

QFX5100-24Q

QFX5100 Switch Specifications

Hardware

Switching Capacity
- QFX5100-48S: 1.44 Tbps/1.07 Bpps
- QFX5100-24Q: 2.56 Tbps/1.44 Bpps
- QFX5100-96S: 2.56 Tbps/1.44 Bpps

Weight
- QFX5100-48S: 21.8 lb (9.9 kg)
- QFX5100-24Q: 21.6 lb (9.8 kg)
- QFX5100-96S: 32.5 lb (14.74 Kg)

Dimensions (HxWxD)
- QFX5100-48S: 1.72 x17.36 x 20.48 in (4.37 x 44.09 x 52.02 cm)
- QFX5100-24Q: 1.72 x17.36 x 20.48 in (4.37 x 44.09 x 52.02 cm)
- QFX5100-96S: 3.46 x 17.36 x 22.44 in (8.8 x 44.09 x 57 cm)
- Switching mode: Cut-through and store-and-forward
- Front-to-back or back-to-front airflow (for hot aisle/cold aisle deployment)
- Management and rear console port connections

Power Consumption[§]
- QFX5100-48S-AFO/AFI: 150 W
- QFX5100-24Q-AFO/AFI: 161 W

[§]Measured with 10GbE passive DAC and QSFP passive DAC

Predicted Mean Time Between Failures (MTBF)
- QFX5100-48S-AFO/AFI: 200,000 hours
- QFX5100-24Q-AFO/AFI: 200,000 hours

FIT Rate
- QFX5100-48S-AFO/AFI: 4,987
- QFX5100-24Q-AFO/AFI: 4,987

Interface Options
- QFX5100-48S-AFO/AFI:
  - 1GbE SFP: 48 (24 copper 1GbE)
  - 10GbE SFP+: 48/72 (with breakout cable)
  - 40GbE QSFP+: 6
- QFX5100-24Q-AFO/AFI:
  - 1GbE SFP: N/A
  - 10GbE SFP+: 96/104 (with breakout cable)
  - 40GbE QSFP+: 24/32 (with 2 x QFX-EM-4Q)

- QFX5100-96S-AFO/AFI:
  - 1GbE SFP: 96 (48 Copper 1GbE)
  - 10GbE SFP+: 104 (with breakout cable)
  - 40GbE QSFP+: 8
- Each QSFP+ port can be configured as a 4 x 10GbE interface
- Each QSFP+ port can be configured as a 40 Gbps port
- USB port
- Console port
- 2 management ports: 1 RJ-45 and 1 SFP
- Supported transceiver and direct attach cable
- SFP+ 10GbE optical modules
- SFP+ DAC cables: 1/3/5 m twinax copper and 1/3/5/7/10 m active twinax copper
- SFP GbE optical and copper module
- QSFP+ to SFP+ 10GbE direct attach break-out copper (1/3 m twinax copper cable)

Rack Installation Kit
- Versatile four post mounting options for 19-in server rack or datacom rack

Airflow
- Front-to-back and back-to-front cooling
- Redundant variable-speed fans to reduce power draw

Power Supply and Fan Modules
- Dual redundant (1+1) and hot-pluggable power supplies
- 110-240 V single phase AC power
- -36 to -72 V DC power
- Redundant (N+1) and hot-pluggable fan modules for front-to-back and back-to-front airflow

Performance Scale (Unidimensional)
- MAC addresses per system: 288,000*
- VLAN IDs: 4,091
- Number of link aggregation groups (LAGs): 128
- Number of ports per LAG: 32
- FCoE Scale:
  - Number of FCoE VLANs/FC Virtual Fabrics: 4,095
- Firewall filters: 4000
- IPv4 unicast routes: 128,000 prefixes; 208,000 host routes; 64** ECMP paths

* MAC address table uses a hash-based scheme to program entries; therefore, some entries may not be programmed due to hash index collision.
** Roadmap

Specifications (continued)

Performance Scale (Unidimensional) (continued)
- IPv4 multicast routes: 104,000
- IPv6 multicast routes: 52,000
- IPv6 unicast routes: 64,000 prefixes
- ARP entries: 48,000
- Jumbo frame: 9,216 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
- Layer 2 protocol tunneling (L2PT)
- Private VLAN (PVLAN)*
- QinQ*
- VLAN Translation
- Static MAC address assignment for interface
- Per VLAN MAC learning (limit)
- MAC learning disable
- Layer 2 protocol tunneling (L2PT)
- Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)

Link Aggregation
- Multi-chassis link aggregation (MC-LAG)
- Redundant Trunk Group (RTG)
- LAG load sharing algorithm—bridged or routed (unicast or multicast) traffic:
  - IP: 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
  - FCoE packet: Source ID (SID), Destination ID (DID), Originator exchange ID (OXID), source port

Layer 3 Features (IPv4)
- Static routing
- Routing policy
- Routing protocols (RIP, OSPF, IS-IS, BGP)
- Virtual Router Redundancy Protocol (VRRP)
- Bidirectional Forwarding Detection (BFD) protocol
- Virtual router
- Dynamic Host Configuration Protocol (DHCP) relay
- Proxy Address Resolution Protocol (ARP)

Multicast Features
- Internet Group Management Protocol (IGMP): v1, v2, V3
- IGMP snooping: v1, v2, v3
- IGMP Filter
- PIM-SM
- 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
- IP Source Guard
- Dynamic ARP Inspection (DAI)
- Sticky MAC address
- 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
- 12 hardware queues per port (8 unicast and 4 multicast)
- Strict priority queuing (LLQ), smoothed deficit weighted round-robin (SDWRR), weighted random early detection (WRED), weighted tail drop
- 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

Data Center Bridging (DCB)
- Priority-based flow control (PFC)—IEEE 802.1Qbb
- Enhanced Transmission Selection (ETS)—IEEE 802.1Qaz
- 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)
- Virtual Fiber Channel Gateway (VFG)
- FCoE session path learning
- FCoE session health monitoring
- Graceful restart for FIP snooping
- FC-BB-6 VN2VN snooping

*Roadmap

Specifications (continued)

Virtual Chassis*
- 40GbE and 10GbE as Virtual Chassis port
- Virtual Chassis Routing Engine (RE) election
- Virtual Chassis pre-provisioning (plug and play)
- Auto-LAG formation of Virtual Chassis ports
- Mixed Virtual Chassis support
- FCoE transit across Virtual Chassis members
- QoS on Virtual Chassis ports
- Local Designated Forwarding
- Unified ISSU
- Graceful RE switchover (GRES)
- Nonstop routing (NSR)
- Nonstop bridging (NSB)
- Monitor distributed aggregate interface
- Control plane protection for virtual RE

High Availability
- Topology-independent in-service software upgrade (TISSU)
- Bidirectional Forwarding Detection (BFD)
- Uplink failure detection (UFD)

MPLS
- 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
- IPv6 tunneling (6PE) (via IPv4 MPLS backbone)
- MPLS Operation, Administration, and Maintenance (OAM)—LSP ping
- IPv4 L3 VPN (RFC 2547, 4364)

Server Virtualization Management
- Junos Space Virtual Control
- IEEE 802.1Qbg (VEPA hairpin forwarding)
- VMware NSX VXLAN L2 Gateway*
- VXLAN OVSDB*

Management and Operations
- 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
- sFlow v5
- 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

Traffic Mirroring
- Port-based
- LAG port
- VLAN-based
- Filter-based
- Mirror to local
- Mirror to remote destinations (L2 over VLAN)

Standards Compliance

IEEE Standard
- IEEE standard
- 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
- IEEE 802.1Qau*
- IEEE 802.1Qbg*

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

*Roadmap

Specifications (continued)

Supported RFCs (continued)

- 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 1155 SMI
- RFC 1157 SNMPv1
- RFC 1212, RFC 1213, RFC 1215 MIB-II, Ethernet-Like MIB and TRAPs
- RFC 1850 OSPFv2 MIB
- RFC 1901 Introduction to Community-based SNMPv2
- RFC 2011 SNMPv2 for Internet protocol using SMIv2
- RFC 2012 SNMPv2 for transmission control protocol using SMIv2
- RFC 2013 SNMPv2 for user datagram protocol using SMIv2
- RFC 2233, The Interfaces Group MIB Using SMIv2
- RFC 2287 System Application Packages MIB
- RFC 2570 Introduction to Version 3 of the Internet-standard Network Management Framework
- RFC 2571 An Architecture for describing SNMP Management Frameworks (read-only access)
- RFC 2572 Message Processing and Dispatching for the SNMP (read-only access)
- RFC 2576 Coexistence between SNMP Version 1, Version 2, and Version 3
- RFC 2578 SNMP Structure of Management Information MIB
- RFC 2579 SNMP Textual Conventions for SMIv2
- RFC 2580 Conformance Statements for SMIv2
- RFC 2665 Ethernet-like interface MIB
- RFC 2787 VRRP MIB
- RFC 2790 Host Resources MIB
- RFC 2819 RMON MIB
- RFC 2863 Interface Group MIB
- RFC 2932 IPv4 Multicast MIB
- RFC 3410 Introduction and Applicability Statements for Internet Standard Management Framework
- RFC 3411 An architecture for describing SNMP Management Frameworks
- RFC 3412 Message Processing and Dispatching for the SNMP
- RFC 3413 Simple Network Management Protocol (SNMP)— (all MIBs are supported except the Proxy MIB)
- RFC 3414 User-based Security Model (USM) for version 3 of 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

Approvals

Safety

- CAN/CSA-C22.2 No. 60950-1 (2007) Information Technology Equipment—Safety
- UL 60950-1 (2nd Ed.) Information Technology Equipment—Safety
- EN 60950-1 (2005) Information Technology Equipment—Safety
- IEC 60950-1 (2005) Information Technology Equipment—Safety (All country deviations): CB Scheme report.
- EN 60825-1 +A1+A2 (1994) Safety of Laser Products—Part 1: Equipment Classification
- GR-63-Core (2006) Network Equipment, Building Systems (NEBS) Physical Protection
- GR-1089-Core (2006) EMC and Electrical Safety for Network Telecommunications Equipment
- SR-3580 (1995) NEBS Criteria Levels (Level 3)

EMC

- FCC 47CFR, Part 15 Class A (2009) USA Radiated Emissions
- EN 55022 Class A (2006)+ A1 2007 European Radiated Emissions
- VCCI Class A (2007) Japanese Radiated Emissions
- BSMI CNS 13438 and NCC C6357 Taiwan Radiated Emissions
- AS/NZS CISPR22:2009

Specifications (continued)

Environmental Compliance

RoHS Compliant — Restriction of Hazardous Substances (ROHS) 6/6

(50) — China Restriction of Hazardous Substances (ROHS)

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

WEEE — Waste Electronics and Electrical Equipment (WEEE)

♺ — Recycled material

80 PLUS SILVER — 80 Plus Silver PSU Efficiency

Telco
· Common Language Equipment Identifier (CLEI) code

Environmental Ranges
· Operating temperature: 32º to 104º F (0º to 40º C)
· Storage temperature: -40º to 158º F (-40º to 70º C)
· Operating altitude: up to 2,000
· Relative humidity operating: 5% to 90% (noncondensing)
· Relative humidity non-operating: 0% to 95% (noncondensing)

Juniper Networks Service and Support

Juniper Networks leads the market in performance-enabling services designed to accelerate, extend, and optimize your high-performance QFabric System. Our services enable you to maximize operational efficiency, reduce costs, and minimize risk while achieving a faster time-to-value for your network.

By leveraging best practices from across the industry, you get the maximum levels of QFabric System performance, designed and delivered by the world's leading professional QFabric technology experts.

For more information, please visit www.juniper.net/us/en/ products-services.

Ordering Information

Model Number	Description
Switch Hardware	
QFX5100-48S-AFI	QFX5100, 48 SFP+/SFP ports, 6 QSFP ports, redundant fans, redundant AC power supplies, back-to-front airflow
QFX5100-48S-AFO	QFX5100, 48 SFP+/SFP ports, 6 QSFP ports, redundant fans, redundant AC power supplies, front-to-back airflow
QFX5100-48S-DC-AFI	QFX5100, 48 SFP+/SFP ports, 6 QSFP ports, redundant fans, redundant DC power supplies, back-to-front airflow
QFX5100-48S-DC-AFO	QFX5100, 48 SFP+/SFP ports, 6 QSFP ports, redundant fans, redundant DC power supplies, front-to-back airflow
QFX5100-24Q-AFI	QFX5100, 24 QSFP ports, 2 expansion slots, redundant fans, redundant AC power supplies, back-to-front airflow
QFX5100-24Q-AFO	QFX5100, 24 QSFP ports, 2 expansion slots, redundant fans, redundant power AC supplies. front-to-back airflow
QFX5100-24Q-DC-AFI	QFX5100, 24 QSFP ports, 2 expansion slots, redundant fans, redundant DC power supplies, back-to-front airflow
QFX5100-24Q-DC-AFO	QFX5100, 24 QSFP ports, 2 expansion slots, redundant fans, redundant DC power supplies, front-to-back airflow
QFX5100-96S-AFI	QFX5100, 96 SFP/SFP+,8 QSFP+, redundant fans, redundant AC power supplies, back-to-front airflow
QFX5100-96S-AFO	QFX5100, 96 SFP/SFP+,8 QSFP+, redundant fans, redundant AC power supplies, front-to-back airflow
QFX5100-96S-DC-AFI	QFX5100, 96 SFP/SFP+,8 QSFP+, redundant fans, redundant DC power supplies, back-to-front airflow
QFX5100-96S-DC-AFO	QFX5100, 96 SFP/SFP+, 8 QSFP+, redundant fans, redundant DC power supplies, front-to-back airflow
Modules and Spares	
QFX-EM-4Q	4 port QSFP+ expansion module for QFX5100-24Q
JPSU-650-AC-AFO	AC 650W PSU, front-to-back airflow for QFX5100-48S, QFX5100-24Q
JPSU-650-AC-AFI	AC 650W PSU, back-to-front airflow for QFX5100-48S, QFX5100-24Q
JPSU-650-DC-AFO	DC 650W PSU, front-to-back airflow for QFX5100-48S, QFX5100-24Q
JPSU-650-DC-AFI	DC 650W PSU, back-to-front airflow for QFX5100-48S, QFX5100-24Q
JPSU-850-AC-AFO	AC 850W PSU, front-to-back airflow for QFX5100-96S
JPSU-850-AC-AFI	AC 850W PSU, back-to-front airflow for QFX5100-96S
JPSU-850-DC-AFO	DC 850W PSU, front-to-back airflow for QFX5100-96S
JPSU-850-DC-AFI	DC 850W PSU, back-to-front airflow for QFX5100-96S
QFX5100-FAN-AFI	Back-to-front Airflow Fan Module for QFX5100-96S
QFX5100-FAN-AFO	Front-to-back Airflow Fan Module for QFX5100
QFX5100-EM-BLNK	Blank module for QFX5100
QFX5100-96S-FANAFO	Front-to-back Airflow Fan Module for QFX5100-96S
QFX5100-96S-FANAFI	Back-to-front Airflow Fan Module for QFX5100-96S

Ordering Information (continued)

Optics and Transceivers

Model Number	Description
QFX-SFP-10GE-SR	SFP+ 10GBASE-SR 10 Gigabit Ethernet Optics, 850 nm for up to 300 m transmission on multimode fiber (MMF)
QFX-SFP-10GE-USR	SFP+ 10 Gigabit Ethernet Ultra Short Reach Optics, 850 nm for 10 m on OM1, 20 m on OM2, 100 m on OM3 multimode fiber (MMF)
QFX-SFP-10GE-LR	SFP+ 10GBASE-LR 10 Gigabit Ethernet Optics, 1,310 nm for 10 km transmission on single mode fiber-optic (SMF)
QFX-SFP-10GE-ER	SFP+ 10GBASE-ER 10 Gigabit Ethernet Optics, 1,550 nm for 40 km transmission on single-mode fiber (SMF)
EX-SFP-10GE-ZR	SFP+ 10GBASE-ZR 10 Gigabit Ethernet Optics, 1,550 nm for 80 km transmission on single-mode fiber (SMF)
QFX-SFP-DAC-1M	SFP+ 10 Gigabit Ethernet Direct Attach Copper (twinax copper cable) 1 m
QFX-SFP-DAC-3M	SFP+ 10 Gigabit Ethernet Direct Attach Copper (twinax copper cable) 3 m
QFX-SFP-DAC-5M	SFP+ 10 Gigabit Ethernet Direct Attach Copper (twinax copper cable) 5 m
QFX-SFP-DAC-1MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active twinax copper cable) 1 m
QFX-SFP-DAC-3MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active twinax copper cable) 3 m
QFX-SFP-DAC-5MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active twinax copper cable) 5 m
QFX-SFP-DAC-7MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active twinax copper cable) 7 m
QFX-SFP-DAC-10MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active twinax copper cable) 10 m
QFX-QSFP-DAC-1M	QSFP+ to QSFP+ Ethernet Direct Attach Copper (twinax copper cable) 1m passive
QFX-QSFP-DAC-3M	QSFP+ to QSFP+ Ethernet Direct Attach Copper (twinax copper cable) 3m passive
JNP-QSFP-DAC-5M	QSFP+ to QSFP+ Ethernet Direct Attach Copper (twinax copper cable) 5m passive
QFX-QSFP-DACBO-1M	QSFP+ to SFP+ 10 Gigabit Ethernet Direct Attach Breakout Copper (twinax copper cable) 1m
QFX-QSFP-DACBO-3M	QSFP+ to SFP+ 10 Gigabit Ethernet Direct Attach Breakout Copper (twinax copper cable) 3m
QFX-QSFP-40G-SR4	QSFP+ 40GBASE-SR4 40 Gigabit Optics, 850 nm for up to 150 m transmission on MMF
QFX-QSFP-40G-ESR4	QSFP+ 40GBASE-ESR4 40 Gigabit Optics, 300 m(400 m) with OM3(OM4) MMF
JNP-QSFP-40G-LR4	QSFP+ 40GBASE-LR4 40 Gigabit Optics, 1310nm for up to 10km Transmission on SMF
QFX-QSFP-DACBO-1M	QSFP+ to SFP+ 10 Gigabit Ethernet Direct Attach Break out Copper (twinax copper cable) 1 m
QFX-QSFP-DACBO-3M	QSFP+ to SFP+ 10 Gigabit Ethernet Direct Attach Break out Copper (twinax copper cable) 3 m
QFX-SFP-1GE-T	SFP 1000BASE-T Copper Transceiver Module for up to 100 m transmission on Cat5
QFX-SFP-1GE-SX	SFP 1000BASE-SX Gigabit Ethernet Optics, 850 nm for up to 550 m transmission on MMF
QFX-SFP-1GE-LX	SFP 1000BASE-LX Gigabit Ethernet Optics, 1,310 nm for 10 km transmission on SMF

Advanced Feature Licenses

Model Number	Description
QFX-JSL-EDGE-ADV1	QFX3500, QFX3600 and QFX5100-48S Advanced Feature License for IS-IS, BGP, and MPLS
QFX5100-HDNSE-LIC	QFX5100-24Q and QFX5100-96S Advanced Feature License for IS-IS, BGP and MPLS
QFX-VCF-LIC	License for Virtual Chassis Fabric

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at **www.juniper.net.**

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at +1-866-298-6428 or authorized reseller.