



Product Overview

The Juniper Networks QFX3500 Switch delivers a high-performance, ultra low latency, feature rich Layer 2 and Layer 3 solution for the most demanding data center environments. Featuring standards-based Fibre Channel I/O and iSCSI convergence capabilities in a compact form factor, the QFX3500 is a versatile, high density 10GbE platform that also delivers a fabric-ready edge solution for the Juniper Networks QFabric architecture.

Product Description

The high-performance Juniper Networks® QFX3500 Switch addresses a wide range of deployment scenarios, which include traditional data centers, virtualized data centers, high-performance computing, network-attached and iSCSI storage, FCoE convergence, and cloud computing.

Featuring 48 dual-mode small form-factor pluggable transceiver (SFP+/SFP) ports and four quad small form-factor pluggable plus (QSFP+) ports in a 1 U form factor, the QFX3500 Switch delivers feature rich Layer 2 and Layer 3 connectivity to networked devices such as rack servers, blade servers, storage systems, and other switches in highly demanding, high-performance data center environments. For end-to-end convergence, the QFX3500 has extensive DCB capabilities including specific iSCSI support. For FCoE-based converged server edge access environments, the QFX3500 is also a standards-based Fibre Channel over Ethernet (FCoE) Transit Switch and FCoE to Fibre Channel (FCoE-FC) Gateway, enabling customers to protect their investments in existing data center aggregation and Fibre Channel storage area network (SAN) infrastructures. The Juniper Networks QFX3500 is a highly flexible data center switch that offers several deployment options including Fabric, Virtual Chassis and MC-LAG.

When deployed with other components of the Juniper Networks QFabric™ family, including the QFX3100 QFabric Director and the QFX3600-I QFabric Interconnect (in a QFX3000-M QFabric System) or QFX3008-I QFabric Interconnect (in a QFX3000-G QFabric System), the QFX3500 delivers a fabric-ready edge solution that contributes to a high-performance, low latency, single-tier data center fabric.

By converting any 10GbE or 40GbE port into a Virtual Chassis connection, the QFX3500 can leverage all the benefits of Juniper's industry-leading Virtual Chassis technology including simplified management and operations, as well as high availability. As a leaf node in Juniper's new Virtual Chassis Fabric architecture, the QFX3500 switch offers deterministic rack-to-rack throughput, high availability and a smooth transition path from 1GbE to 10/40GbE. The QFX3500 can be deployed along with with Juniper QFX3600, EX4300 and QFX5100 switches in the same Virtual Chassis and Virtual Chassis Fabric configuration.

For added configuration flexibility, up to 36 of the QFX3500's 48 pluggable SFP+ ports can be used in 10GbE or 1GbE mode with up to 18 of the 1GbE ports being copper. The remaining 12 ports can be used to support 2, 4, or 8 Gbps Fibre Channel modes as well as 10GbE. Each of the four QSFP+ high speed ports can operate in 4x10GbE mode and are capable of supporting 40 Gbps optics* in the future, providing investment protection.

*Roadmap




Figure 1: QFX3500 Switch (rear view)

Junos OS

The QFX3500 high-performance switch runs Junos OS, Juniper’s powerful and robust network operating system. Key Junos OS features that enhance the functionality and capabilities of the QFX3500 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 QFX3500 switches can be managed by 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.

* Refer to the specifications section for features shipping now vs. those available with a future software release.

Virtual Chassis Technology

The QFX3500 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 QFX3500 switches can participate in the same Virtual Chassis configuration with any combination of EX4300, QFX3600, and QFX5100 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 QFX5100, QFX3500, QFX3600 or EX4300 switches as leaf nodes. This architecture provides any-rack-to-any-rack deterministic throughput and low latency, while significantly simplifying network operations through a single point of management. The QFX3500, EX4300, QFX3600 and QFX5100 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 QFX3500 investments.

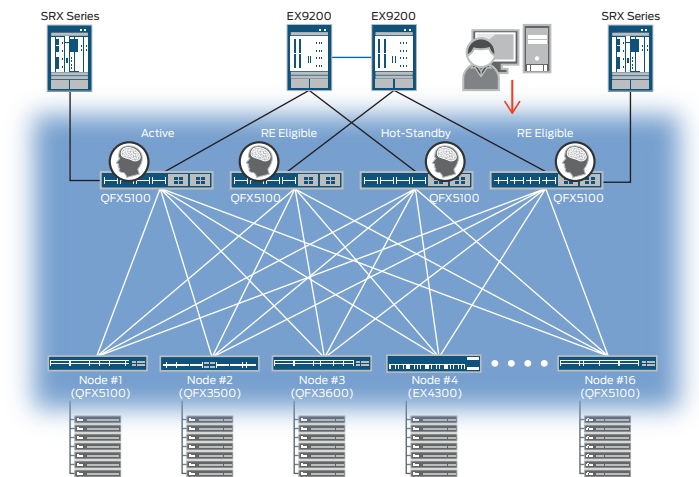


Figure 2: Virtual Chassis Fabric data center deployment with a mix of 1GbE and 10GbE

QFabric System

The QFX3500 can also act as a QFabric Node device in a QFX3000-M and QFX3000-G QFabric System. With a mixture of QFX3500, QFX3600, and QFX5100 Nodes, a QFabric System can scale on a node by node basis.

Data Center Deployments

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

The QFX3500 Switch delivers the ultra low latency, lossless high density 10GbE interfaces, and FCoE Transit Switch and FCoE-FC Gateway functionality demanded by today’s data center—all in a compact form factor designed to support high-performance, converged data center access networks. The QFX3500’s low power consumption optimizes the switch’s power use effectiveness (PUE) ratio to reduce data center operating costs, while front-to-back air flow meets hot and cold aisle isolation requirements.

The QFX3500 is an ideal platform for:

- High-performance Layer 2 and Layer 3 access for enterprise, HPC, cloud and Web-tier data centers (Figure 2 and Figure 3)
- High-performance data center bridging (DCB), storage and I/O convergence environments (Figure 4):
 - FCoE Transit Switch
 - FCoE-FC Gateway
 - iSCSI Transit Switch
- ULL Peering Router between co-location/ exchange (Figure 5)
- L3 Aggregation Switch (Figure 6)
- Fabric edge deployments in the Juniper Networks QFabric architecture (Figure 5)

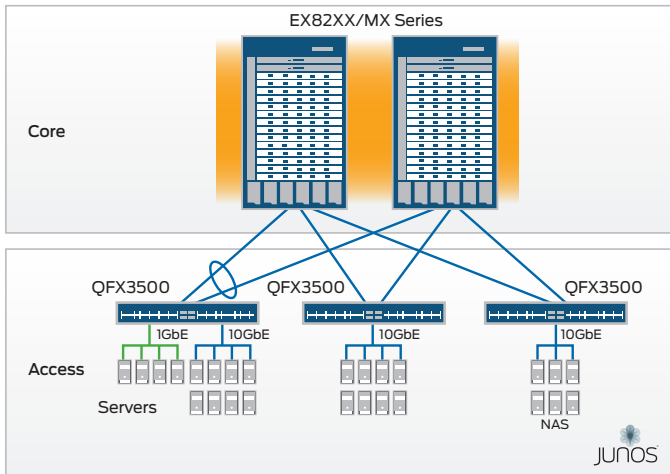


Figure 3: QFX3500 high-performance Ethernet L2 and L3 access deployment scenario

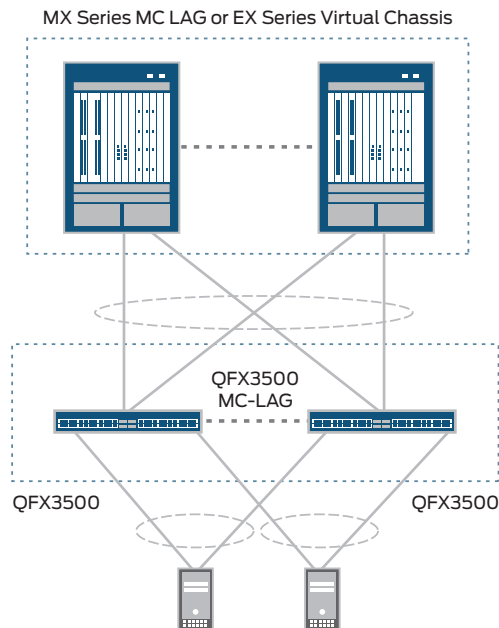


Figure 4: QFX3500 as a Layer 2 access switch with server dual homing using MC-LAG technology

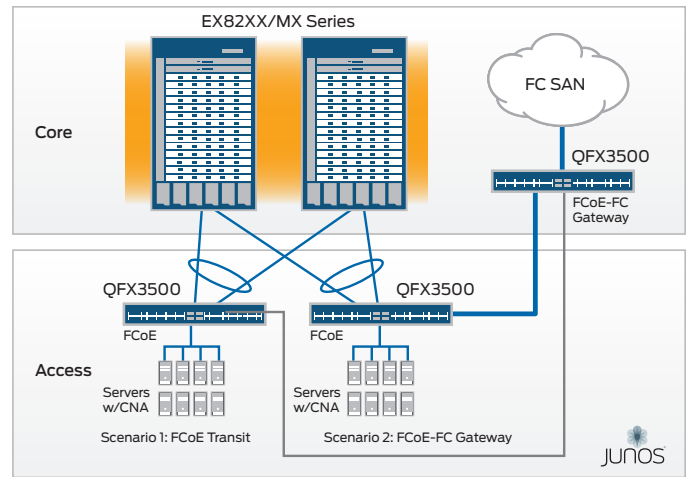


Figure 5: High-performance DCB, storage, and I/O convergence deployment scenarios

High-Performance L2 / L3 Access Deployments

With its compact, 1U form factor, high density 10GbE ports, and rear-port design, the QFX3500 is ideal for top-of-rack deployments in high-performance, highly demanding data centers. Featuring 1.28 Tbps throughput, 960 Mpps forwarding capacity, and low power consumption, the QFX3500 is a highly efficient solution for aggregating high-performance servers in the data center. The QFX3500 can operate in both cut-through and store-and-forward modes and delivers sustained wire-speed switching with sub-microsecond latency and low jitter for any packet size (including jumbo frames) in both modes. The QFX3500 supports extensive Layer 2 features which enable the device to support high-density 10GbE/FCoE Layer 2 access deployments. With 32-way ECMP support for all routing protocols (RIP, IS-IS, OSPF and BGP), it is also possible to enable Layer 3 access from a QFX3500 top-of-rack switch to an upstream aggregation switch.

The QFX3500 also supports active-active server dual homing using multi-chassis link aggregation (MC-LAG) technology for maximum resiliency and full bisectional bandwidth while eliminating the need for Spanning Tree Protocol (STP). MC-LAG enables a client device to form a logical LAG interface using two QFX3500 switches, providing redundancy and load balancing between the QFX3500 devices, multi-homing support, and a loop-free Layer 2 network without running STP.

At this time, MC-LAGs support only Layer 2 features. A client device on one end of an MC-LAG with one or more physical links in a LAG does not need to be aware of the MC-LAG; they only need to support the LACP standard. At the other end of the MC-LAG are two QFX3500 switches, each of which has one or more physical links connected to a single client device. The QFX3500 devices coordinate with each other to ensure that data traffic is forwarded properly. Both unicast and multicast traffic are supported with QFX3500 MC-LAG technology.

High-Performance DCB, Storage, and I/O Convergence Deployments

The QFX3500 is a fully IEEE DCB- and T11 FC-BB-5-based FCoE Transit Switch and FCoE-FC Gateway, delivering a high-performance solution for converged server edge access environments. The QFX3500 provides configurable ports capable of 1GbE, 10GbE, and 2/4/8 Gbps FC connectivity.

- FCoE Transit Switch:** As an FCoE Transit Switch, the QFX3500 provides a pure IEEE DCB converged access layer between FCoE-enabled servers and an FCoE-enabled Fibre Channel SAN (see Figure 4, Scenario 1). The QFX3500 offers a full featured DCB implementation that provides strong monitoring capabilities on the top-of-rack switch for SAN and LAN administration teams, while maintaining a clear separation of management. In addition, FC Initiation Protocol (FIP) snooping provides perimeter protection, ensuring that the presence of an Ethernet layer does not impact existing SAN security policies. The FCoE Transit Switch functionality, including Priority-based Flow Control (PFC), Enhanced Transmission Selection (ETS), and Data Center Bridging Exchange (DCBX) with the FCoE Application TLV, are included as part of the default software; no additional licenses are required.
- FCoE-FC Gateway:** In FCoE-FC Gateway mode, the QFX3500 eliminates the need for FCoE enablement in the SAN backbone. Organizations can add a converged access layer and interoperate with existing SANs without disrupting the network. The QFX3500 allows up to 12 ports to be converted to Fibre Channel without requiring additional switch hardware modules, and gateway functionality can be soft provisioned with a software license to protect existing investments. The QFX3500 acts as an N_Port Virtualizer using N_Port ID Virtualization (NPIV) proxy functionality between FCoE-enabled servers and traditional Fibre Channel SANs. As a top-of-rack switch with FCoE-FC Gateway functionality, the QFX3500 presents itself as an FCoE-enabled switch to the rack or blade servers, and as a group of logical FC servers to the traditional Fibre Channel SAN (see Figure 4, Scenario 2).
- iSCSI Transit Switch:** As an iSCSI Transit Switch, the QFX3500 provides a pure IEEE DCB-converged network between iSCSI-enabled servers and iSCSI-enabled storage. The QFX3500 offers a full-featured DCB implementation that provides strong monitoring capabilities on the top-of-rack switch for storage and LAN administration teams while maintaining a clear separation of management. The iSCSI Transit Switch functionality, notably Priority-based Flow Control (PFC), Enhanced Transmission Selection (ETS), and Data Center Bridging Exchange (DCBX) including the iSCSI Application TLV, are included as part of the default software; no additional licenses are required.

Low Latency Peering Router Between CoLo and Exchange

With sub-microsecond latency, high 10GbE port densities in a compact 1U form factor, and L3 unicast and multicast routing features such as BGP, OSPF and PIM, the QFX3500 delivers a highly effective switch/router solution that can be deployed at co-location facilities where market feed data is received from financial exchanges.

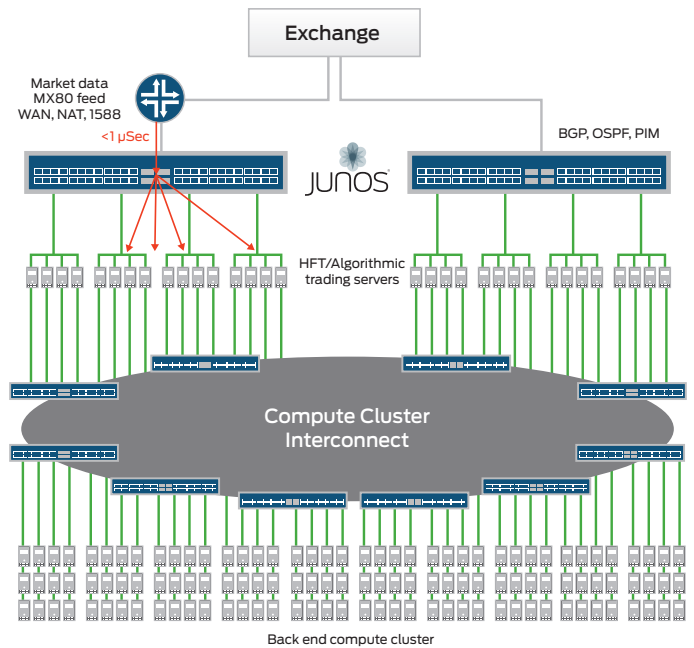


Figure 6: ULL Peering Router between CoLo and Exchange

L3 Aggregation Switch

The QFX3500 supports L3 unicast and multicast routing protocols as well as the Junos suite of routing features, making it suitable for deployment as a small-scale L3 aggregation switch to support small 10GbE SaaS clusters as well as small virtualized enterprise data centers. In these scenarios, L3 routing starts at the top-of-rack switch. The QFX3500 also supports virtual routers, allowing administrators to divide the switch into multiple independent instances, each with their own routing tables. This enables customers to be aggregated and their traffic isolated to segment the network without using multiple devices. Customer-specific routing instances are bound to the customer-owned interfaces.

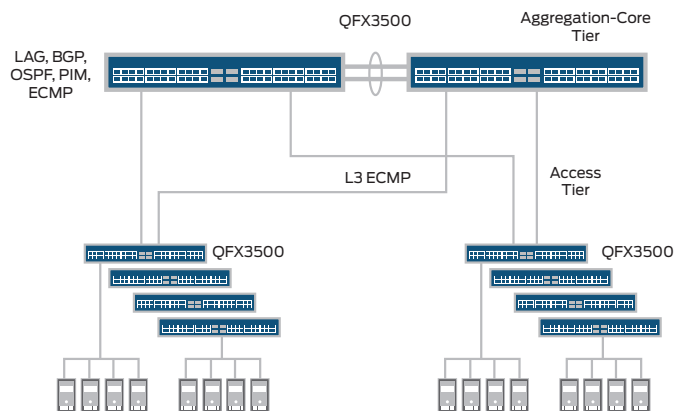


Figure 7: L3 Aggregation Switch

Fabric Edge Solution of Juniper Networks QFabric Architecture

The QFX3500 also delivers a fabric-ready solution that, with a simple configuration change, easily transitions into a QFabric Node edge device that serves as a building block for a Juniper Networks QFabric System, protecting existing investments while providing a seamless architectural migration to a single-tier data center network. The QFX3500 can be used as a QFabric Node in either a QFX3000-M or QFX3000-G system.

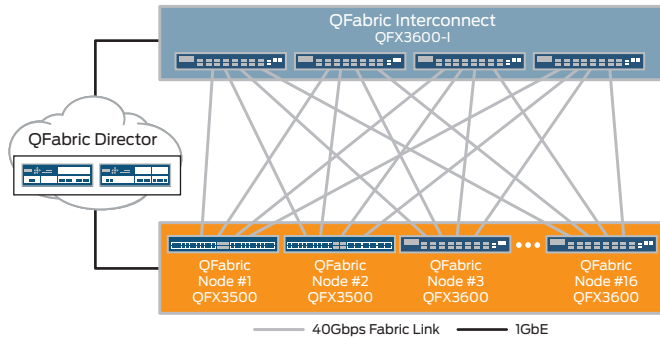


Figure 8: QFX3500 as a QFabric Node Link in a QFX3000-M QFabric System

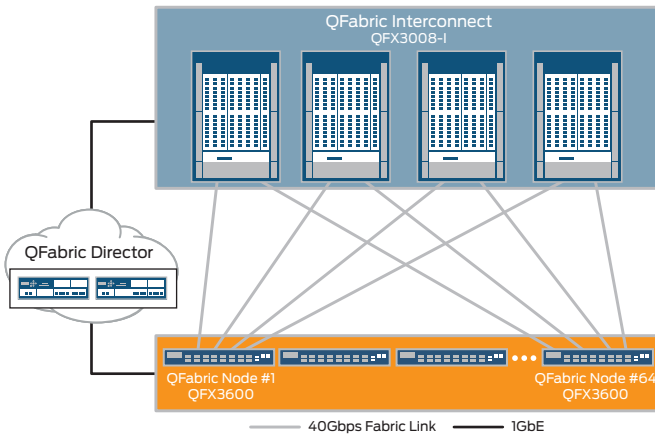


Figure 9: QFX3500 as a QFabric Node in a QFX3000-G QFabric System

Management Capabilities

The following system management options are available for the QFX3500 Switch.

- The standard Junos OS command-line interface (CLI) and SNMP module offers the same granular management capabilities and scripting parameters found in any router or switch powered by Junos OS.
- QFX3500 fault, performance, and configuration data can also be exported to leading third-party management systems such as HP OpenView, IBM Tivoli, and Computer Associates Unicenter software, providing a complete, consolidated view of network operations.
- The QFX3500 is supported by Juniper Networks Junos® Space, an open, programmable application platform for hosting a comprehensive suite of network operational application tools that provide an open approach for automating the deployment and operation of a Juniper infrastructure. Explicitly designed to allow partners and customers to build and deploy smart, simple,

and easy-to-use applications, Junos Space provides multiple management and infrastructure applications for managing Juniper resources and assets, including inventory management, device and interface configuration, automated software management and deployment, and event-driven fault management. These Junos Space applications offer predefined automation schemes and best practice templates to enable rapid and accurate deployments.

- The QFX3500 also supports the Junos XML management protocol and Junos XML management protocol automation tools. The Junos XML management protocol facilitates the quick and easy development of custom applications and scripts using the XML programming language. Junos XML management protocol automation tools provide early detection and automatic resolution of potential problems related to the operating system.

Features and Benefits

- **Performance:** The QFX3500 delivers 1.28 Tbps of throughput and 960 Mpps of switching capacity to sustain wire-speed switching with ultra low latency and low jitter. All ports can run at full wire-speed capacity with full performance in both L2 and L3 mode, with the option to operate in either cut-through or store-and-forward mode. BFD support for routing protocols ensures fast convergence and provides failure detection on any kind of path between switches, including direct and indirect Layer 3 links.
- **Rich QoS:** The QFX3500 is designed for general-purpose 10GbE as well as iSCSI, NAS, and FC/FCoE deployments. The QFX3500 supports rich QoS functionality in hardware, including multiple dedicated hardware queues for unicast and multicast traffic, a sophisticated two-tier scheduler, and rich QoS functionality delivered via Junos software, which includes classification, marking, policing and shaping functionality. A dynamic buffer (9MB) allocation scheme with both dedicated and pooled buffer spaces allows the switch to react to congestion and short periods of micro-bursts across ports and priorities while supporting fairness lossless forwarding
- **Reliability:** The QFX3500 is designed with robust high availability features that include redundant AC and DC power supplies and fan modules to ensure hardware availability. Control plane and data plane separation, combined with the Junos OS high availability design, ensures maximum systems-level availability.
- **Server Virtualization:** The QFX3500 Switch supports a large media access control (MAC) address table that enables large-scale server virtualization deployment. The QFX3500 is also supported by Junos Space Virtual Control, which unifies physical and virtual infrastructures to provide network operators with a comprehensive view into the complete end-to-end network. In addition, the QFX3500 Switch supports the IEEE 802.1Qbg standard, hairpin forwarding, to enable Virtual Ethernet Port Aggregator (VEPA) deployment.
- **Automation:** The QFX3500 switches support a number of features for network automation and plug-and-play operations. Features include zero-touch provisioning, operations and event scripts, and atomic rollback. The switch also offers support for integration with VMware vCenter, Puppet, and OpenStack.
- **Intelligent Buffer Management:** The QFX3500 switches have a total of 8MB shared buffers. While 25% of the total buffer space is dedicated, the rest is shared among all ports and is user configurable. The intelligent buffer mechanism in the QFX3500 effectively absorbs traffic bursts while providing deterministic performance, significantly increasing performance over static allocation.

- **Insight Technology for Analytics:** The QFX3500 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:** QFX3500 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 QFX3500 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 QFX3500 provides an IEEE data center bridging (DCB) converged network between FCoE-enabled servers and an FCoE-enabled Fibre Channel storage area network (SAN). The QFX3500 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.
The QFX3500 offers a full-featured FCoE to FC gateway capability.
- **Energy Efficiency and Environmental Friendliness:** The QFX3500 is an environmentally conscious green solution that lowers operational expenses. The switch consumes less than 5 watts per 10GbE port, while variable speed fans dynamically adjust their speed based on ambient temperature to optimize operating power. With maximum power consumption of 365 W and nominal power of 230 W for AC-powered systems and 250 W for DC-powered systems, the QFX3500 is certified for Silver PSU Efficiency at 85%. The QFX3500 is also certified for environmentally responsible compliance with labels such as Reduction of Hazardous Substances (ROHS), Registration, Evaluation, Authorisation and Restriction of Chemicals (REACH), and Waste Electronics and Electrical Equipment (WEEE).



*Roadmap

QFX3500 Switch Specifications

Hardware

System

- Switching capacity: 1.28 Tbps and 960 Mpps switching capacity
- Switching mode: Cut-through and store-and-forward
- Weight: 30.75 lb (13.95 kg) with dual power supplies and fan modules
- Front-to-back or back-to-front airflow (for hot aisle/cold aisle deployment)
- Rear port connections

Interface Options

- 48 SFP+ ports
- 36 dual-mode 10GbE or GbE ports
- 12 dual-mode 10GbE or 2/4/8 Gbps FC ports
- Four QSFP+ ports
- Each QSFP+ port can be configured as a 4 x 10GbE interface
- Each QSFP+ port can be configured as a 40 Gbps port*
- Dual management ports (copper and fiber versions)
- USB port
- Console port

Supported Transceiver and Direct Attach Cable

- SFP+ 10GbE optical modules
- SFP+ DAC cables: 1/3/5/7 m twinax copper and 1/3/5/7/10 m active twinax copper
- SFP Fibre Channel 2/4/8 Gbps module
- SFP GbE optical and copper module
- QSFP+ to SFP+ 10GbE direct attach break-out copper (1/3 m twinax copper cable)

Dimension

- Height: 1.75 in (4.45 cm), 1 U
- Width: 17.25 in (43.82 cm)
- Depth: 28 in (71.12 cm)

Rack Installation Kit

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

Display and Ports

- LCD panel for system monitoring

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 V to -72 V DC power
- Redundant and hot-pluggable fan modules for front-to-back and back-to-front airflow

Performance Scale

- MAC addresses per system: 120,000
- VLAN IDs: 4,095
- Number of link aggregation groups (LAGs): 63
- Number of ports per LAG: 32
- FCoE Scale:
 - Number of FCoE VLANs / FC Virtual Fabrics: 4,095
 - Number of VFGs: 12 (up to 4 VLANs per virtual gateway)
 - Number of FCoE VF ports: 376 in untrusted mode; 2,500 in trusted mode
- Firewall filters: 1,500

QFX3500 Switch Specifications (continued)

Performance Scale (continued)

- IPv4 unicast routes: 16,000 prefixes; 8,000 host routes; 32 ECMP paths
- Multicast groups: 3,500
- 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)
- Private VLAN (PVLAN)
- VLAN Translation
- Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
- Jumbo frame (9,216 bytes)

Link Aggregation

- Multi-chassis link aggregation (MC-LAG)
- 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, ISIS, BGP)
- Virtual Router Redundancy Protocol (VRRP)
- Bidirectional Forwarding Detection (BFD) protocol
- Virtual router

- Dynamic Host Configuration Protocol (DHCP) relay
- Proxy ARP

Multicast Features

- Internet Group Management Protocol (IGMP): v1, v2, v3
- IGMP snooping: v1, v2, v3
- IGMP Filter
- PIM-SM
- Multicast Source Directory 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 Address Resolution Protocol (ARP) support
- Storm control, port error disable, and auto-recovery
- Control plane denial of service (DoS) protection
- 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 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 TLVs
- Ethernet Congestion Management (QCN) – IEEE 802.1Qau*

Fibre Channel over Ethernet (FCoE)

- FCoE Transit Switch (FIP Snooping ACL Installation)
- FCoE-FC Gateway (with FC feature license)
- Virtual Fiber Channel Gateway (VFG)
- FCoE Session Path Learning
- FCoE Session Health Monitoring
- Graceful restart for FIP Snooping
- FC-BB-6 VN2VN Snooping

*Roadmap

QFX3500 Switch Specifications (continued)

Fibre Channel Standard

- FC port speeds: 2, 4, 8 Gbps
- FC port types: Fabric mode N_Port and VF_Port
- FC classes of service: Class 3
- FC services: N_Port Virtualization and N_Port ID Virtualization
- FCoE support: FC-BB-5 FC-BB_E, including FIP Snooping
- Configurable FCoE-FC load balancing

High Availability

- Bi-directional Forwarding Detection (BFD)
- Uplink failure detection
- Graceful Restart for routing protocols and FC-FCoE Gateway

Virtual Chassis

- 40GbE and 10GbE as Virtual Chassis port
- Virtual Chassis RE election
- Virtual Chassis pre-provisioning (plug and play)
- Auto-LAG formation of Virtual Chassis ports
- FCoE transit across Virtual Chassis members
- QoS on Virtual Chassis ports
- Local Designated Forwarding*
- NSSU
- Graceful Restart (GRES)
- Non-stop Routing (NSR)*
- Non-stop Bridging (NSB)*

MPLS

- Static LSPs
- RSVP based signaling of LSPs
- LDP based signaling of LSPs
- LDP Tunneling (LDP over RSVP)
- MPLS CoS
- MPLS ACL/Policers
- MPLS LSR Support
- IPv6 Tunneling (6PE) (via IPv4 MPLS backbone)
- MPLS OAM - LSP ping
- IPv4 L3 VPN (RFC 2547, 4364)

Server Virtualization Management

- Junos Space Virtual Control
- IEEE 802.1Qbg (VEPA hairpin forwarding)

Management and Operations

- Role-based CLI management and access
- CLI via console, telnet, or SSH
- Show and debug commands, statistics
- Extended ping and traceroute
- Junos OS configuration rescue and rollback
- Image rollback
- SNMP v1/v2/v3
- Junos XML management protocol
- sFlow v5
- J-Web
- Zero Touch Provisioning (ZTP)
- OpenStack Neutron Plus-in
- Puppet

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 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 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 1812 Requirements for IP Version 4 routers
- RFC 2030 SNTP, Simple Network Time Protocol
- RFC 2068 HTTP server
- RFC 2138 RADIUS Authentication
- RFC 2139 RADIUS Accounting
- RFC 2267 Network ingress filtering
- RFC 2453 RIP v2
- RFC 2131 BOOTP/DHCP relay agent and Dynamic Host Configuration Protocol (DHCP) server
- RFC 1591 Domain Name System (DNS)
- RFC 2338 VRRP

*Roadmap

QFX3500 Switch Specifications (continued)

Supported RFCs (continued)

- RFC 2328 OSPF v2 (edge mode)
- RFC 1587 OSPF not-so-stubby area (NSSA) Option
- RFC 2154 OSPF w/Digital Signatures (Password, MD-5)
- RFC 2370 OSPF Opaque link-state advertisement (LSA) Option
- RFC 3623 Graceful OSPF Restart
- RFC 4812 OSPF Restart Signaling
- RFC 1745 BGP4/IDRP for IP—OSPF Interaction
- RFC 1772 Application of the Border Gateway Protocol in the Internet
- RFC 1997 BGP Communities Attribute
- RFC 2385 Protection of BGP Sessions via the TCP MD5 Signature Option
- RFC 2439 BGP Route Flap Damping
- 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 3392 Capabilities Advertisement with BGP-4
- RFC 4271 A 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 4893 BGP Support for Four-octet AS Number Space
- RFC 5396 Textual Representation of Autonomous System (AS) Numbers
- RFC 5668 4-Octet AS Specific BGP Extended Community
- RFC 2362 PIM-SM (edge mode)
- RFC 1112 IGMP v1
- RFC 2236 IGMP v2
- RFC 3376 IGMP v3 (source-specific multicast include mode only)
- RFC 3446, Anycast RP
- RFC 3569 SSM
- RFC 3618 MSDP
- RFC 4724 Graceful Restart Mechanism for BGP
- RFC 5880 Bidirectional Forwarding Detection (BFD)
- RFC 5176 Dynamic Authorization Extensions to RADIUS
- RFC 2474 Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers
- RFC 2597 Assured Forwarding PHB Group
- RFC 2598 An Expedited Forwarding PHB
- RFC 2698 A Two Rate Three Color Marker
- RFC 2697 A Single Rate Three Color Marker

Supported MIBs

- RFC 1155 SMI
- RFC 1157 SNMPv1
- RFC 1212, RFC 1213, RFC 1215 MIB-II, Ethernet-Like MIB and TRAPs
- 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 2790 Host Resources MIB
- RFC 2819 RMON MIB
- RFC 2863 Interface Group 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
- Supported MIBs (continued)
- 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
- RFC 2932 IPv4 Multicast MIB
- RFC 2787 VRRP MIB
- RFC 1850 OSPFv2 MIB

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) (pending)

QFX3500 Switch Specifications (continued)

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

Environmental Compliance



Restriction of Hazardous Substances (ROHS) 6/6



Silver PSU Efficiency



Recycled material



Waste Electronics and Electrical Equipment (WEEE)



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



China Restriction of Hazardous Substances (ROHS)

Telco

- Common Language Equipment Identifier (CLEI) code

Environmental Ranges

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

Juniper Networks Services 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 performance, designed and delivered by the world's leading professional QFabric service experts.

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

Ordering Information

Model Number	Description
Base Hardware	
QFX3500-48S4Q-ACR	QFX3500, 48 SFP+/SFP and 4 QSFP ports, redundant dual AC power supply, front to back air flow
QFX3500-48S4Q-ACRB	QFX3500, 48 SFP+/SFP and 4 QSFP ports, redundant dual AC power supply, port side-to-FRU side air flow and RJ-45 management port
QFX3500-48S4Q-ACR-F	QFX3500, 48 SFP+/SFP and 4 QSFP ports, redundant dual AC power supply, FRU side-to-port side air flow and fiber management port
QFX3500-48S4Q-ACRB-F	QFX3500, 48 SFP+/SFP and 4 QSFP ports, redundant dual AC power supply, port side-to-FRU side air flow and fiber management port
QFX3500-RB-ACRB	QFX3500, 48 SFP+/SFP and 4 QSFP ports, redundant dual AC power supply, port side-to-FRU side air flow, RJ-45 management port and removable mounting brackets
QFX3500-RB-ACRB-F	QFX3500, 48 SFP+/SFP and 4 QSFP ports, redundant dual AC power supply, port side-to-FRU side air flow, fiber management port and removable mounting brackets
Hardware Spares	
QFX3500-48S4Q-AFI	QFX3500, 48 SFP+/SFP and 4 QSFP ports, redundant fan trays, FRU side-to-port side air flow (management module and power supplies are required and sold separately)
QFX3500-48S4Q-AFO	QFX3500, 48 SFP+/SFP and 4 QSFP ports (spare), redundant fan trays, port side-to-FRU side air flow (management module and power supplies are required and sold separately)
QFX3500-RB-AFO	QFX3500, 48 SFP+/SFP and 4 QSFP ports with removable mounting brackets, accessory kit, dual fan trays, port side-to-FRU side airflow (management module and power supplies are required and sold separately)
QFX3500-RB-RMK	Adjustable 4-post rack-mount kit for QFX3500-RB-AFO
QFX3500-48S4Q	QFX3500, 48 SFP+/SFP and 4 QSFP ports (spare)
QFX3500-MB	Management board for QFX3500-48S4Q (spare)
QFX3500-FANAI	Fan module A (Intake) for QFX3500-48S4Q (spare)
QFXC01-PWRACI-650A	650 W power supply A (with Intake air flow) for QFX3500 (spare)
QFX3500-RMAR	Rear rack mount bracket A for QFX3500 (spare)
JPSU-650W-AC-AFI	650 W AC power supply for QFX3500 (PSU-side airflow intake)
JPSU-650W-AC-AFO	650 W AC power supply for QFX3500 (PSU-side airflow exhaust)
JPSU-650W-DC-AFI	650 W DC power supply for QFX3500 (PSU-side airflow intake)
JPSU-650W-DC-AFO	650 W DC power supply for QFX3500 (PSU-side airflow exhaust)
QFX3500-FAN-AFO	Fan module (port side-to-FRU side air flow) for QFX3500-48S4Q (spare)

Ordering Information (continued)

Model Number	Description
Hardware Spares (continued)	
QFX3500-MB-RJ45-AFO	Management board with 1GbE RJ-45 interface for QFX3500-48S4Q (port side-to-FRU side air flow) (spare)
QFX3500-MB-SFP-AFI	Management board with SFP interface for QFX3500-48S4Q (FRU side-to-port side air flow) (spare)
QFX3500-MB-SFP-AFO	Management board with SFP interface for QFX3500-48S4Q (port side-to-FRU side air flow) (spare)

Power Cables

CBL-EX-PWR-C13-US	AC power cable, US (13 A/125 V, 2.5 meter)
CBL-EX-PWR-C13-EU	AC power cable, Europe (10 A/250 V, 2.5 meter)
CBL-EX-PWR-C13-AU	AC power cable, Australia (10 A/250 V, 2.5 meter)
CBL-EX-PWR-C13-CH	AC power cable, China (10 A/250 V, 2.5 meter)
CBL-EX-PWR-C13-IT	AC power cable, Italy (10 A/250 V, 2.5 meter)
CBL-EX-PWR-C13-JP	AC power cable, Japan (12 A/125 V, 2.5 meter)
CBL-EX-PWR-C13-KR	AC power cable, Korea (10 A/250 V, 2.5 meter)
CBL-EX-PWR-C13-SZ	AC power cable, Switzerland (10 A/250 V, 2.5 meter)
CBL-EX-PWR-C13-UK	AC power cable, UK (10 A/250 V, 2.5 meter)

Optics and Transceivers

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, 1310 nm for 10 km transmission on single mode fiber-optic (SMF)
QFX-SFP-10GE-ER	SFP+ 10GBASE-ER 10 Gigabit Ethernet Optics, 1550 nm for 40 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

Model Number	Description
Optics and Transceivers (continued)	
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-SFP-8GFC-SW	SFP 2/4/8-Gbps Fibre Channel SW Optics
QFX-QSFP-40G-SR4	QSFP+ 40GBASE-SR4 40 Gigabit Optics, 850nm for up to 150m transmission on MMF
QFX-QSFP-40G-ESR4	QSFP+ 40GBASE-ESR4 40 Gigabit Optics, 300m(400m) with OM3(OM4) MMF
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-IGE-T	SFP 1000BASE-T Copper Transceiver Module for up to 100 m transmission on Cat5
QFX-SFP-IGE-SX	SFP 1000BASE-SX Gigabit Ethernet Optics, 850 nm for up to 550 m transmission on MMF
QFX-SFP-IGE-LX	SFP 1000BASE-LX Gigabit Ethernet Optics, 1310 nm for 10 km transmission on SMF

Advanced Feature Licenses

QFX-JSL-EDGE-FC	QFX Series edge feature license for FCoE-FC gateway
QFX-JSL-EDGE-ADVI	QFX Series edge Advanced Feature License for IS-IS, BGP and MPLS Routing

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.

Corporate and Sales Headquarters

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or +1.408.745.2000
Fax: +1.408.745.2100
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
Fax: +31.0.207.125.701

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

Copyright 2014 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos and QFabric are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.