


QFX10000 Modular Ethernet Switches



Product Overview

The QFX10000 line of modular data center spine and core Ethernet switches delivers industry-leading scale, flexibility and openness, with a design that enables the seamless transition from 10GbE and 40GbE interface speeds to 100GbE and beyond. These high-performance, forward-looking switches are designed to help cloud and data center operators extract maximum value and intelligence from their network infrastructure well into the future.

Product Description

The Juniper Networks® QFX10000 line of modular Ethernet switches delivers up to 96 Tbps of system throughput, scalable to over 200 Tbps in the future, to meet the rapid and ongoing traffic growth in the data center. Industry-leading scale and density on the QFX10000 modular switches redefine per-slot economics, enabling customers to do more with less while simplifying network design and reducing OpEx. Based on purpose-built Q5 ASICs, the QFX10000 line delivers unparalleled intelligence and analytics, providing deeper insights into application performance.

The QFX10000 line can be deployed in a number of different network designs and fabrics, including Layer 3 fabric, Junos Fusion, and Juniper MC-LAG for Layer 2 and Layer 3 networks, giving customers complete architectural flexibility. Additionally, the open architecture ensures that customers can innovate on top of Juniper Networks Junos® operating system to accelerate the pace of innovation. The QFX10000 line also supports Junos Fusion, a simple, easy-to-deploy fabric that scales to support mid-to-large data center deployments while simplifying network management and configuration.

The QFX10000 is also a key component of Juniper Networks MetaFabric™ Architecture, which provides a simple, open and smart approach to building data center networks. Additionally, the QFX10000 supports Data Center Interconnect (DCI), critical to workload mobility and application availability.

For networks evolving to SDN, the QFX10000 integrates with VMware NSX SDN controllers and can act as a Virtual Extensible LAN (VXLAN) Layer 2 and Layer 3 gateway. The open and standards-based QFX10000 switches also interoperate with the Open vSwitch Database (OVSDB) to support automated management and control capabilities. The QFX10000 line integrates with OpenContrail SDN controller to allow users to choose their preferred SDN systems.

Two QFX10000 modular chassis options are available, providing full deployment flexibility:

- QFX10008 Ethernet Switch, an 8-slot, 13 U chassis that supports up to eight line cards
- QFX10016 Ethernet Switch, a 16-slot, 21 U chassis that supports up to 16 line cards

Both QFX10000 modular chassis can accommodate any combination of the following QFX10000 Ethernet line cards:

- QFX10000-36Q, a 36-port 40GbE quad small form-factor pluggable plus transceiver (QSFP+) or 12-port 100GbE QSFP28 line card
- QFX10000-30C, a 30-port 100GbE QSFP28/40GbE QSFP+ line card
- QFX10000-60S-6Q, a 60-port 1GbE/10GbE SFP/SFP+ line card with six-port 40GbE QSFP+ / two-port 100GbE QSFP28

Fully configured, a single QFX10016 chassis can support up to 480 100GbE ports, delivering the industry's highest line-rate 100GbE port densities in its class. The QFX10000 switch fabric is capable of delivering up to 7.2 Tbps (full duplex) per slot, and the midplane-less orthogonal interconnect architecture ensures system longevity.

QFX10000 Modular Switches Highlights

- Industry-leading line-rate 100GbE port density with up to 480 100GbE ports in a single chassis
- Up to 96 Tbps Layer 2 and Layer 3 performance, scalable to over 200 Tbps in the future
- Unparalleled investment protection with high density 10GbE, 40GbE, and 100GbE; 400GbE ready
- System longevity with midplane-less orthogonal interconnect architecture
- Highest logical Layer 2 / Layer 3 scale; up to 1M MACs, 2 million host routes, 2 million FIB
- Deep buffers with up to 50ms delay bandwidth buffer
- No head-of-line blocking with virtual output Queue (VoQ)-based architecture
- Flexible network architectures including Layer 3 fabric, Junos Fusion, and Juniper MC-LAG for Layer 2 and Layer 3 networks
- Scalable, plug-and-play Ethernet fabric with Junos Fusion
- Juniper Virtualized Open Network Operating System framework for programmability through APIs
- High availability with Topology-Independent In-Service Software Upgrade (TISSU)
- Next-generation analytics with Cloud Analytics Engine
- Advanced Junos OS features such as BGP add-path, VXLAN routing, MPLS, and FCoE
- Rich automation capabilities with operations and event scripts, Python, Chef and Puppet

Architecture and Key Components

The QFX10000 modular data center spine and core Ethernet switches share a number of architectural elements. The Control Boards (CBs) employed by these switches run Junos OS, which processes all Layer 2 and Layer 3 protocols, while the Switch Fabric modules manage the chassis and provide switching functionality for data traffic coming from line cards.

The QFX10000 line cards, which are common across all modular QFX10000 platforms, include Packet Forwarding Engines (PFEs) that process network traffic, as well as a line-card processor that provides scalable local control.

The QFX10000 Virtual Output Queue (VOQ)-based architecture is designed to scale to very large deployments, with no head-of-line blocking, a single-tier low-latency switch fabric, efficient multicast replication handling, and deep buffering to ensure performance at scale. The horizontal line cards in the front of the chassis directly connect with the vertical switch fabric cards in

the rear of the chassis via orthogonal interconnects, without the need for a midplane. This midplane-less architecture ensures a smooth upgrade process to higher-speed switch fabric cards in the future, providing unparalleled investment protection. The system is designed from the ground up to support 400GbE in future.

To maintain uninterrupted operation, the QFX10000 fan trays cool the line cards, Control Board, and Switch Fabric modules with redundant, variable-speed fans. In addition, the QFX10000 power supplies convert building power to the internal voltage required by the system.

All QFX10000 components are hot-swappable, and all central functions are available in redundant configurations, providing high operational availability by allowing continuous system operation during maintenance or repairs.

QFX10000 Line Cards

The QFX10000 line cards support an extensive set of Layer 2 and Layer 3 services that can be deployed in any combination of L2-L3 applications.

Each QFX10000 line card is unique in its ability to support tri-speed 10GbE, 40GbE and 100GbE connections, enabling customers to transition seamlessly from one speed to another as their needs change. The unprecedented port density, high logical table sizes and up to 50ms deep packet buffers allow customers to deploy the most scalable systems necessary to meet the exponential data growth in the data center.

Each QFX10000 line card is built upon the Juniper Q5 silicon, which supports a wide range of Layer 2 and Layer 3 Ethernet functionality including 802.1Q VLAN, VXLAN, link aggregation, Virtual Router Redundancy Protocol (VRRP), L2 to L3 mapping, and port monitoring. Additionally, the line cards support filtering, sampling, load balancing, rate limiting, class of service (CoS), MPLS, Fibre Channel over Ethernet (FCoE) transit functionality, and other key features needed to deploy a dependable, lossless, high-performance Ethernet infrastructure.

QFX10000 Control Board

The QFX10000 Control Board supports control and management plane functionality with an integrated Routing Engine (RE) that features a quad-core, 2.5 GHz Intel processor with 32 gigabytes of SDRAM and an on-board solid-state drive (SSD) providing 32 GB of storage for Junos OS images and logs. The CB also has a front SSD slot to support a hot swappable SSD for secondary images and other external storage. The CB features AUX, console, and Ethernet ports on the front panel to support out-of-band system management and monitoring, while an external USB port provides a removable media interface to install Junos OS images manually.

The CB's central CPU performs all system control functions and maintains hardware forwarding table and routing protocol states for the QFX10000 modular switches. Dedicated hardware on the CB module supports chassis management functions such as environmental monitoring, while communication between CB modules and individual line cards takes place over a dedicated internal 10GbE out-of-band control interface.

QFX10000 Switch Fabric

The QFX10000 Switch Fabric modules are hot-swappable and serve as the central non-blocking matrix through which all network data passes.

Both the QFX10008 and QFX10016 systems have six Switch Fabric cards that provide N+1 redundancy. The Switch Fabric cards deliver 7.2 Tbps throughput per slot. The switch cards can be upgraded to provide more than 200 Tbps of system throughput in the future.

Switch Fabric modules perform the following key functions:

- Monitor and control system functions
- Interconnect all line cards
- Handle clock and system resets

Power

The QFX10008 contains six power supply bays while the QFX10016 has ten power supply bays to provide complete flexibility for provisioning and redundancy. Each power supply has its own internal fan for cooling. All QFX10000 chassis support both AC and DC power supplies; however, AC and DC supplies cannot be mixed in the same chassis.

The AC supplies on the QFX10008 chassis accept 200 to 240 volts alternating current (VAC) input and deliver 2,700 watts of power to the chassis, while the DC power supplies accept -40 to -72 volts direct current (VDC) input and deliver 2,500 watts of power to the chassis. Each AC and DC power supply has two inputs for feed redundancy.

Cooling

The QFX10000 modular chassis support front-to-back cooling with air being pulled in through the perforations on the Control Boards and the line cards in the front of the chassis. Hot air is exhausted through the fan trays that are placed in front of the fabric cards in the rear of the chassis as well as through the power supplies.

Each chassis hosts two fan tray subsystems with redundant fans. Each fan tray subsystem consists of a fan tray controller and a fan tray. The fan tray controller provides power to the fan trays and manages the fans in the fan tray.

Flexible Network Architectures

The QFX10000 switches support a diverse set of deployment options, including Layer 3 fabric, Junos Fusion, and Layer 2 and Layer 3 networks with Juniper multichassis link aggregation group (MC-LAG). Customers can choose the architecture that best suits their deployment needs and easily adapt and evolve as requirements change over time. The QFX10000 switches serve as the universal building block for all of these 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 non-blocking and predictable performance and scale characteristics. As an example, a two-tier fabric with the QFX10000 switches as the spine and QFX5100 switches as leafs can scale to support more than 50,000 10GbE server ports at 3:1 oversubscription.

One of the most complicated tasks in building an IP fabric is assigning all of the details such as IP addresses, BGP AS numbers, routing policy, loopback addresses, and many other implementation details. Automating the creation of an IP fabric at a large scale is also difficult. Juniper has created the OpenClos project to provide free, open-source tools to automate the creation of IP fabrics in the data center. OpenClos is 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 network switch configuration files and even a cabling plan.

Junos Fusion: Junos Fusion, an IEEE 802.1Br standards-based Ethernet fabric, is ideal for mid-to-large data center customers who want simplicity at scale. With QFX10000 switches deployed as the “aggregation” devices in a Junos Fusion configuration, customers can reduce network complexity and operational costs by collapsing access layer network elements called “satellite devices” into a single, logical point of control and management. Junos Fusion also increases deployment flexibility by centralizing management and automating device configuration while reducing total cost of ownership (TCO).

The open, easy-to-deploy fabric provides a single point of management and plug-and-play operation for networks requiring 6,000 10GbE server ports or more, and can scale to support more than 12,000 10GbE server ports in the future. The choice of 1GbE, 10GbE, 40GbE and, in the future, 100GbE “satellite” devices in a Junos Fusion configuration provides a seamless transition for customers as networks evolve. Customers can also choose to operate their network in an intelligent distributed forwarding model for intra-rack traffic, thereby offloading the aggregation devices.

Juniper MC-LAG: The QFX10000 switches also eliminate Spanning Tree Protocol in traditional Layer 2 networks when deployed in a MC-LAG configuration. 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 highest availability for applications.

Overlays: Customers can also deploy overlay networks over Layer 3 fabrics as well as over Junos Fusion and MC-LAG deployments in the future. These overlay networks use VXLAN in the data plane and can support NVGRE in the future. Overlay networks can be built without an SDN controller using EVPN as the control plane technology, or they can be orchestrated through VMware NSX or OpenContrail using OVSDB as the control and management plane technology.

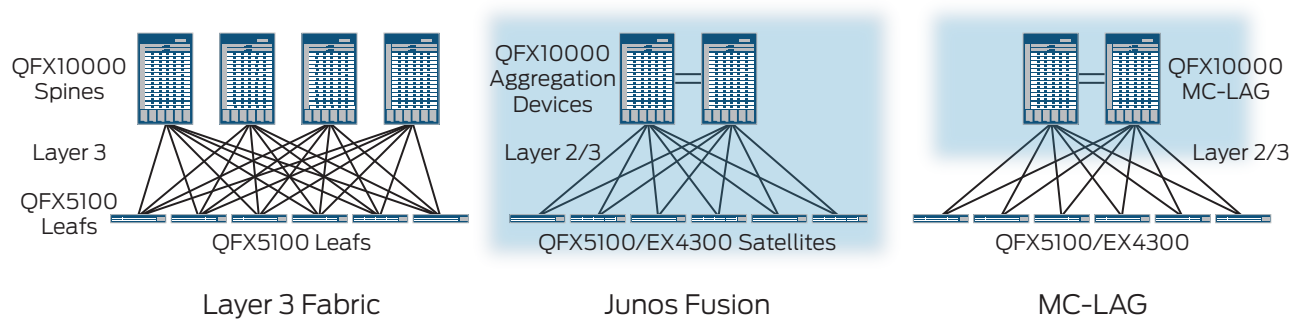


Figure 1: QFX10000 modular switches can be deployed in Layer 3 fabric, Junos Fusion or MC-LAG configurations.

Carrier-Class Operating System

The QFX10000 line of switches runs the same Junos OS used by all other Juniper Networks EX Series and QFX Series Ethernet Switches, as well as the Juniper Networks 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 QFX10000 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

Juniper Virtualized Open Network Operating System Framework

The QFX10000 switches have an open software architecture that enables customers to innovate along with Juniper to accelerate the pace of innovation. Customers can create and run applications alongside Junos; these applications can then follow their own software release schedule.

This new solution offers modularity and direct programmability through APIs. The platform and Packet Forwarding Engine (PFE) modules are separate from the control module(s). Customers can create and run guest applications, either as a VM or a Linux daemon, alongside Junos and communicate with Junos and

also directly with platform and PFE modules through normalized programmable APIs. Having the ability to build applications that can program the QFX10000 control plane, data plane and the platform itself, the customers can implement specific and unique network forwarding functions and automate the control and management by integrating with the customer's SDN controller and orchestration tools. The Juniper QFX10000 programmable APIs have a stable abstraction so that the customer applications can work across Juniper platforms as well as future Junos releases. Based on Apache Thrift, the Juniper QFX10000 supports standardized API definition via Interface Description Language (IDL) while allowing the use of customer's favorite programming language.

Other services such as analytics and automation can also run as separate daemons and can be directly accessed through orchestration tools.

Cloud Analytics Engine

The QFX10000 line of switches supports the new Juniper Networks Cloud Analytics Engine, a next-generation analytics tool designed for complex, dynamic data centers of any size. Using network data analysis to improve application performance and availability, Cloud Analytics Engine performs data collection, correlation and visualization to help customers better understand workload and application behavior across the physical and virtual infrastructure.

Cloud Analytics Engine provides these major capabilities:

- Application visibility and performance management, by controlling application flows and workload placement
- Capacity planning and optimization, by detecting hotspots and monitoring latency and microbursts
- Troubleshooting and root cause analysis, by correlating overlay and underlay networks

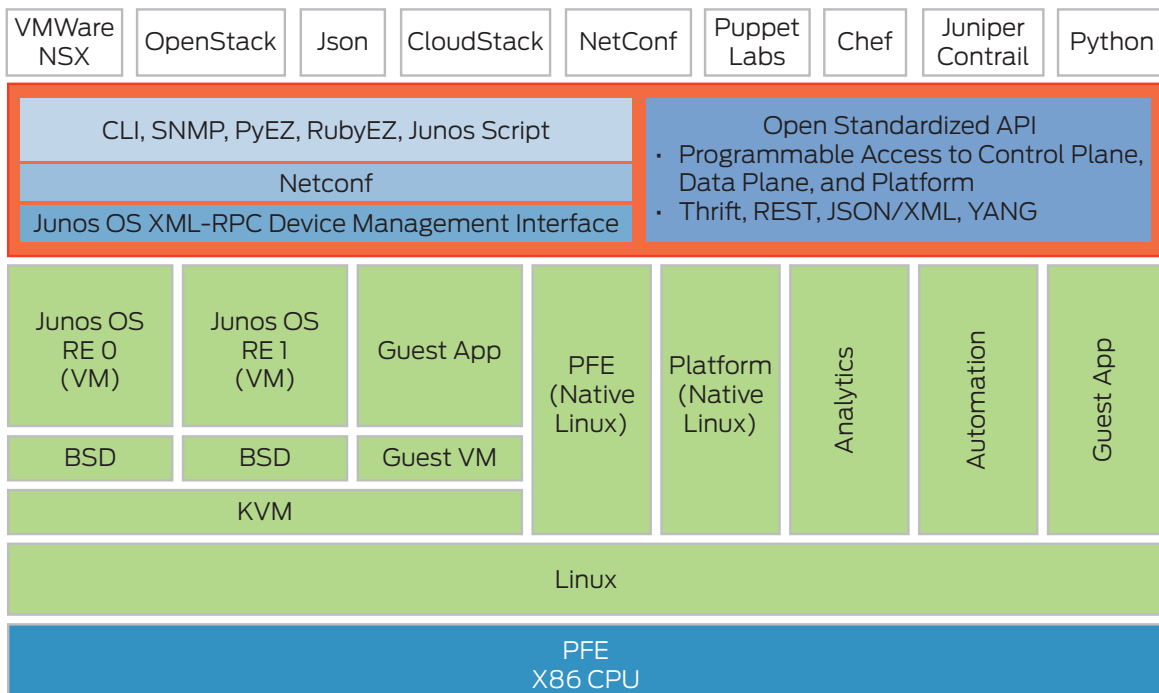


Figure 2: Juniper Virtualized Open Network Operating System Framework

Junos Space Network Director

The QFX10000 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—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.

Features and Benefits

High Availability

QFX10000 modular spine and core switches deliver a number of high availability features that ensure uninterrupted, carrier-class performance. Each QFX10000 chassis includes an extra slot to accommodate a redundant RE module that serves as a backup in hot-standby mode, ready to take over in the event of a master RE failure. If the master fails, the integrated Layer 2 and Layer 3 graceful Routing Engine switchover (GRES) feature of Junos OS, working in conjunction with the nonstop active routing (NSR) and

nonstop bridging (NSB) features, ensures a seamless transfer of control to the backup, maintaining uninterrupted access to applications, services, and IP communications. The QFX10000 modular switches also support topology-independent in-service software upgrade (TISSU) that enables them to seamlessly move to a newer software version while maintaining data plane traffic intact.

Virtual Output Queue (VOQ)

The QFX10000 switches support a virtual output queue (VOQ)-based architecture designed for very large deployments. VOQ refers to a queue on the egress port that is maintained by the ingress PFE. With VOQ architecture, packets are queued and dropped on ingress during congestion with no head-of-line blocking.

Automation

The QFX10000 switches support a number of features for network automation and plug-and-play operations. Features include operations and event scripts, automatic rollback, and Python scripting. The switches also support integration with VMware NSX, OpenContrail, Puppet, OpenStack, and CloudStack.

MPLS

QFX10000 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.

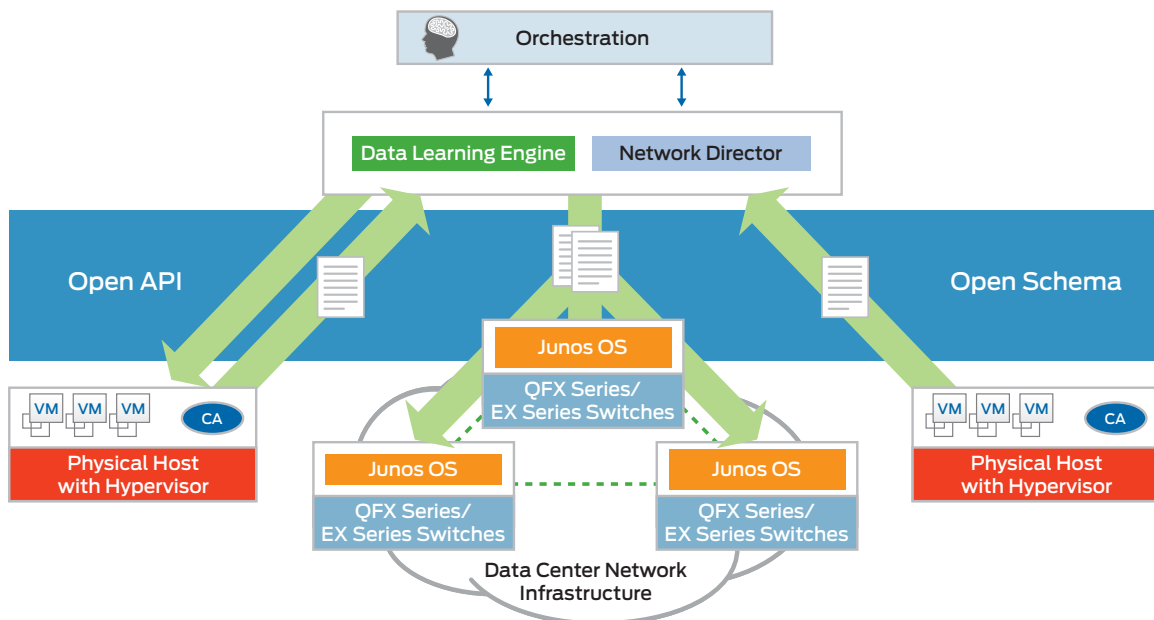


Figure 3: Cloud Analytics Engine components.


VXLAN

The QFX10000 supports Layer 2 and Layer 3 gateway services that enable VXLAN-to-VLAN connectivity at any tier of the data center network, from server access to the edge. The QFX10000 integrates with NSX through data plane (VXLAN) and control and management plane (OVSDB) protocols to centrally automate and orchestrate the data center network.

FCoE

As Fibre Channel over Ethernet (FCoE) transit switches, the QFX10000 line provides an IEEE data center bridging (DCB) converged network between FCoE-enabled servers and an FCoE-enabled Fibre Channel storage area network (SAN). The QFX10000 offers a full-featured DCB implementation that provides strong monitoring capabilities, helping SAN and LAN administration teams maintain clear management separation. 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.



Specifications

Hardware

Table 1: QFX10000 System Capacity

	QFX10008	QFX10016
System throughput	Up to 48 Tbps	Up to 96 Tbps
Forwarding capacity	Up to 16 Bpps	Up to 32 Bpps
Maximum bandwidth/slot	7.2 Tbps/slot	
Maximum 10GbE port density	1,152	2,304
Maximum 40GbE port density	288	576
Maximum 100GbE port density	240	480

Table 2: QFX10000 Chassis Specifications

	QFX10008	QFX10016		QFX10008	QFX10016
Dimensions (W x H x D):	17.4 x 22.55 x 32 in (44.2 x 57.3 x 81.3 cm)	17.4 x 36.65 x 35 in (44.2 x 93 x 88.9 cm)	Total number of I/O slots	8	16
Rack units	13 U	21 U	Fabric	<ul style="list-style-type: none"> Up to 86.4 Tbps raw fabric capacity Up to 7.2 Tbps (full duplex) per slot usable fabric capacity 5+1 fabric redundancy 	<ul style="list-style-type: none"> Up to 172.8 Tbps raw fabric capacity Up to 7.2 Tbps (full duplex) per slot usable fabric capacity 5+1 fabric redundancy
Weight			Power	Holds up to six power supplies: <ul style="list-style-type: none"> -40 to -72 VDC 200 to 240 VAC 	Holds up to ten power supplies: <ul style="list-style-type: none"> -40 to -72 VDC 200 to 240 VAC
Base configuration	280 lbs (127 kg)	491 lbs (223 kg)	Cooling	<ul style="list-style-type: none"> Front-to-back airflow Dual hot-swappable fan trays with redundant fans 	
Redundant configuration	322 lbs (146 kg)	596 lbs (270 kg)	Latency	<ul style="list-style-type: none"> As low as 2.5 microseconds within a PFE As low as 5.5 microseconds across PFEs 	
Fully loaded chassis	446 lbs (202 kg)	844 lbs (383 kg)	Warranty	Juniper standard one-year warranty	
Architecture	<ul style="list-style-type: none"> Midplane-less orthogonal interconnect architecture Separate dedicated data, control, and management planes Secure, modular architecture that isolates faults Transparent failover and network recovery 				
Operating System	Junos OS				
Control Board	<ul style="list-style-type: none"> Master and backup CBs (1+1 redundancy) Intel Quad Core 2.5GHz CPU 32 gigabytes SDRAM 50 gigabytes onboard SSD Pluggable solid-state drive (SSD) USB 2.0 storage interface Four SFP/SFP+ MACsec-capable ports Fiber (SFP) or 10/100/1000BASE-T (RJ-45) Ethernet management port One console port 				




Table 3: Line Card Specifications

	QFX10000-36Q	QFX10000-30C	QFX10000-60S-6Q
Dimensions (W x H x D):	17.2 x 1.89 x 20.54 in (43.7 x 4.8 x 52.2 cm)	17.2 x 1.89 x 20.54 in (43.7 x 4.8 x 52.2 cm)	17.2 x 1.89 x 20.54 in (43.7 x 4.8 x 52.2 cm)
Weight	22.6 lbs (10.2 kg)	27 lbs (12.2 kg)	21.4 lbs (9.7 kg)
Maximum 10GbE port density	144	N/A	84
Maximum 40GbE port density	36	30	6
Maximum 100GbE port density	12	30	2
Buffer	12 GB	12 GB	8 GB

Table 4: QFX10000 System Power Consumption

	QFX10008 Typical Power	QFX10008 Reserved Power	QFX10016 Typical Power	QFX10016 Reserved Power
Base system	1,517 W	2,472 W	3,989 W	6,124 W
Redundant system	1,765 W	2,866 W	4,618 W	7,023 W

Table 5: QFX10000 Component Power Consumption

	Typical Power	Reserved Power
QFX10000-36Q Line Card	520 W	690 W
QFX10000-30C Line Card	890 W	1,100 W
QFX10000-60S-6Q Line Card	365 W	480 W
QFX10008 Switch Fabric	170 W	225 W
QFX10016 Switch Fabric	510 W	675 W
QFX10000 Control Board	50 W	125 W
QFX10008 Fan Tray	225 W	475 W
QFX10016 Fan Tray	475 W	975 W

Software

Table 6: QFX10000 Modular Switches Scale (Uni-dimensional)

Media access control (MAC) addresses	576,000 (roadmap: up to 1 million)
ARP entries	127,000 (roadmap: up to 256,000)
Jumbo frames	9,216 Bytes maximum
VLANs	4,093 (roadmap: up to 16,000)
Forwarding information base (FIB) (IPv4/IPv6)	2 million
Host routes	Up to 2 million
ECMP	64-way
IPv4 Multicast routes	Up to 128,000
IPv6 Multicast routes	Up to 128,000
Multicast groups	Up to 128,000
Filters	Up to 8,000 per PFE
Filter terms	Up to 64,000 per PFE
Policers	Up to 8,000 per PFE
Output queues per port	8
Virtual Output Queues	384,000 per PFE
Link aggregation groups (LAG)	128 (roadmap: 1,000)
Distributed BFD sessions	200/line card
Members / LAG	64
Mirroring sessions	48
GRE tunnels	2,000 (roadmap: 4,096)
MPLS IPv4 Layer 3 VPNs (Advanced Feature License)	4,096

Layer 2 Features

- 802.1D – Spanning Tree Protocol (STP)
- 802.1w – Rapid Spanning Tree Protocol (RSTP)
- 802.1s – Multiple Spanning Tree Protocol (MSTP)
- VLAN Spanning Tree Protocol (VSTP)
- 802.1AB Link Layer Discovery Protocol (LLDP)
- VLAN Registration Protocol
- QinQ

Link Aggregation

- 802.3ad – Link Aggregation Control Protocol (LACP)
- Multi-chassis Link Aggregation (MC-LAG)

Layer 3 Features

- Static routing
- RIP v1/v2
- OSPF v1/v2
- OSPF v3
- Filter-based forwarding
- Virtual Router Redundancy Protocol (VRRP)
- IPv6
- Bidirectional Forwarding Detection (BFD)
- Virtual routers
- Unicast reverse path forwarding (uRPF)
- Loop-free alternate (LFA)
- BGP (Premium Feature license or Advanced Feature license)
- IS-IS (Premium Feature license or Advanced Feature license)
- DHCP v4/v6 relay
- VR-aware DHCP
- IPv4 / IPv6 over GRE tunnels (interface-based)

Multicast

- Internet Group Management Protocol (IGMP) v1/v2/v3
- 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)

Firewall Filters

- Ingress and egress L2-L4 access control lists (ACLs):
 - Port ACLs
 - VLAN ACLs
 - Router ACLs
- Control plane denial-of-service (DoS) protection

Quality of Service (QoS)

- Single rate three color policer
- Two rate three-color policer
- Congestion prevention:
 - Weighted random early detection (WRED)
 - Tail drop
 - Explicit congestion notification (ECN) marking
- Priority-based scheduling:
 - Low-latency queuing with strict high priority
 - Weighted round-robin (WRR) queuing

MPLS (Advanced Feature License)

- LDP
- RSVP
- LDP tunneling (LDP over RSVP)
- Fast reroute (FRR)
- IPv6 tunneling (6PE)
- Ingress, transit, and egress label-switched paths (LSPs)
- BFD

Overlays (Premium Feature License or Advanced Feature License)

- Integration with Juniper Networks Contrail
- Integration with VMware NSX SDN controller
- VXLAN
- Open vSwitch Database (OVSDB)
- VXLAN L2 and L3 gateway
- Ethernet VPN (EVPN)-VXLAN

High Availability

- Graceful Routing Engine switchover (GRES)
- Nonstop active routing (NSR)
- Nonstop bridging (NSB)
- Topology-independent in-service software upgrade (TISSU)*

Timing

- Precision Timing Protocol (PTP)
 - Transparent clock

Visibility

- Switched Port Analyzer (SPAN)
- RSPAN
- ERSPAN
- Firewall filter-based port mirroring
- sFlow v5
- High frequency statistics monitoring
- Cloud Analytics Engine*

Data Center Bridging

- 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)

Management and Operations

- Junos Space Network Director
- Junos OS CLI via console, telnet, SSH
- Out-of-band management: Serial; 10/100/1000BASE-T Ethernet
- Role-based CLI management and access
- Junos XML management protocol
- ASCII configuration file
- SNMP v1/v2/v3
- RADIUS
- TACACS+
- Extensive MIB support
- Junos OS configuration rescue and rollback
- Image rollback
- OpenStack Neutron Plug-in
- Puppet
- Chef
- Python
- Junos OS event, commit and OP scripts
- RMON (RFC 2819): Groups 1, 2, 3, 9
- Network Time Protocol (NTP)
- SSHv2
- Secure copy
- DNS resolver
- System logging
- Environment monitoring
- Temperature sensor
- Configuration backup via FTP/secure copy

Troubleshooting

- Debugging: CLI via console, telnet, or SSH
- Diagnostics: Show, debug, and statistics commands
- Port mirroring
- IP tools: Extended ping and trace
- Juniper Networks commit and rollback

IEEE Compliance

- IEEE 802.1AB: Link Layer Discovery Protocol (LLDP)
- IEEE 802.1ad: QinQ
- IEEE 802.1D-2004: Spanning Tree Protocol (STP)
- IEEE 802.1p: Class-of-service (CoS) prioritization
- IEEE 802.1Q: Virtual Bridged Local Area Networks
- IEEE 802.1s: Multiple Spanning Tree Protocol (MSTP)
- IEEE 802.1w: Rapid Spanning Tree Protocol (RSTP)

* Roadmap

- IEEE 802.3: 10BASE-T
- IEEE 802.3u: 100BASE-T
- IEEE 802.3ab: 1000BASE-T
- IEEE 802.3z: 1000BASE-X
- IEEE 802.3ae: 10-Gigabit Ethernet
- IEEE 802.3ba: 40-Gigabit/100-Gigabit Ethernet
- IEEE 802.3ad: Link Aggregation Control Protocol (LACP)
- IEEE 802.1Qbb: Priority-based Flow Control
- IEEE 802.1Qaz: Enhanced Transmission Selection
- RFC 2154: OSPF with Digital Signatures (password, Message Digest 5)
- RFC 2236: IGMP v2
- RFC 2267: Network Ingress Filtering
- RFC 2270: BGP-4 Dedicated autonomous system (AS) for sites/single provider
- RFC 2283: Multiprotocol Extensions for BGP-4
- RFC 2328: OSPF v2 (Edge mode)
- RFC 2338: VRRP
- RFC 2362: PIM-SM (Edge mode)

RFC Compliance

- RFC 768: UDP
- RFC 783: Trivial File Transfer Protocol (TFTP)
- RFC 791: IP
- RFC 792: Internet Control Message Protocol (ICMP)
- RFC 793: TCP
- RFC 826: ARP
- RFC 854: Telnet client and server
- RFC 894: IP over Ethernet
- RFC 903: Reverse Address Resolution Protocol (RARP)
- RFC 906: TFTP Bootstrap
- RFC 951, 1542: BootP
- RFC 1027: Proxy ARP
- RFC 1058: RIP v1
- RFC 1112: IGMP v1
- RFC 1122: Host Requirements
- RFC 1142: OSI IS-IS Intra-domain Routing Protocol
- RFC 1256: IPv4 ICMP Router Discovery Protocol (IRDP)
- RFC 1492: TACACS+
- RFC 1519: Classless Interdomain Routing (CIDR)
- RFC 1587: OSPF NSSA Option
- RFC 1591: Domain Name System (DNS)
- RFC 1745: BGP4/IDRP for IP-OSPF Interaction
- RFC 1765: OSPF Database Overflow
- RFC 1771: Border Gateway Protocol 4
- RFC 1772: Application of the Border Gateway Protocol in the Internet
- RFC 1812: Requirements for IP Version 4 Routers
- RFC 1965: Autonomous System Confederations for BGP
- RFC 1981: Path maximum transmission unit (MTU) discovery for IPv6
- RFC 1997: BGP Communities Attribute
- RFC 2030: Simple Network Time Protocol (SNTP)
- RFC 2068: HTTP server
- RFC 2080: RIPng for IPv6
- RFC 2081: RIPng Protocol Applicability Statement
- RFC 2131: BOOTP/Dynamic Host Configuration Protocol (DHCP) relay agent and DHCP server
- RFC 2138: RADIUS Authentication
- RFC 2139: RADIUS Accounting
- RFC 2370: OSPF Opaque LSA Option
- RFC 2373: IPv6 Addressing Architecture
- RFC 2375: IPv6 Multicast Address Assignments
- RFC 2385: TCP MD5 Authentication for BGPv4
- RFC 2439: BGP Route Flap Damping
- RFC 2453: RIP v2
- RFC 2460: Internet Protocol, v6 (IPv6) specification
- RFC 2461: Neighbor Discovery for IP Version 6 (IPv6)
- RFC 2462: IPv6 Stateless Address Autoconfiguration
- RFC 2463: ICMPv6
- RFC 2464: Transmission of IPv6 Packets over Ethernet Networks
- RFC 2474: DiffServ Precedence, including 8 queues/port
- RFC 2526: Reserved IPv6 Subnet Anycast Addresses
- RFC 2545: Use of BGP-4 Multiprotocol Extensions for IPv6 Interdomain Routing
- RFC 2547: BGP/MPLS VPNs
- RFC 2597: DiffServ Assured Forwarding (AF)
- RFC 2598: DiffServ Expedited Forwarding (EF)
- RFC 2697: A Single Rate Three Color Marker
- RFC 2698: A Two Rate Three Color Marker
- RFC 2710: Multicast Listener Discovery (MLD) for IPv6
- RFC 2711: IPv6 Router Alert Option
- RFC 2740: OSPF for IPv6
- RFC 2796: BGP Route Reflection (supersedes RFC 1966)
- RFC 2796: Route Reflection
- RFC 2858: Multiprotocol Extensions for BGP-4
- RFC 2893: Transition Mechanisms for IPv6 Hosts and Routers
- RFC 2918: Route Refresh Capability for BGP-4
- RFC 3031: Multiprotocol Label Switching Architecture
- RFC 3032: MPLS Label Stack Encoding
- RFC 3036: LDP Specification
- RFC 3065: Autonomous System Confederations for BGP
- RFC 3176: sFlow
- RFC 3215: LDP State Machine
- RFC 3306: Unicast Prefix-based IPv6 Multicast Addresses
- RFC 3376: IGMP v3
- RFC 3392: Capabilities Advertisement with BGP-4
- RFC 3446: Anycast Rendezvous Point (RP) Mechanism using PIM and MSDP

- RFC 3478: Graceful Restart for Label Distribution Protocol
- RFC 3484: Default Address Selection for IPv6
- RFC 3513: Internet Protocol Version 6 (IPv6) Addressing
- RFC 3569: PIM-SSM PIM Source Specific Multicast
- RFC 3587: IPv6 Global Unicast Address Format
- RFC 3618: Multicast Source Discovery Protocol (MSDP)
- RFC 3623: OSPF Graceful Restart
- RFC 3768: Virtual Router Redundancy Protocol (VRRP)
- RFC 3810: Multicast Listener Discovery Version 2 (MLDv2) for IP
- RFC 3973: PIM-Dense Mode
- RFC 4213: Basic Transition Mechanisms for IPv6 Hosts and Routers
- RFC 4291: IPv6 Addressing Architecture
- RFC 4360: BGP Extended Communities Attribute
- RFC 4364: BGP/MPLS IP Virtual Private Networks (VPNs)
- RFC 4443: ICMPv6 for the IPv6 specification
- RFC 4486: Sub codes for BGP Cease Notification message
- RFC 4552: Authentication/Confidentiality for OSPFv3
- RFC 4604: Using Internet Group Management Protocol Version 3 (IGMPv3)
- RFC 4724: Graceful Restart Mechanism for BGP
- RFC 4798: Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)
- RFC 4861: Neighbor Discovery for IPv6
- RFC 4862: IPv6 Stateless Address Autoconfiguration
- RFC 5095: Deprecation of Type 0 Routing Headers in IPv6
- RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates
- RFC 5306: Restart Signaling for IS-IS
- RFC 5308: Routing IPv6 with IS-IS
- RFC 5340: OSPF for IPv6
- RFC 5880: Bidirectional Forwarding Detection
- RFC 2013: SNMPv2 for user datagram protocol using SMIv2
- RFC 2096: IPv4 Forwarding Table MIB
- RFC 2287: System Application Packages MIB
- RFC 2465: Management Information Base for IP Version 6
- RFC 2570–2575: SNMPv3, user-based security, encryption, and authentication
- 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 2665: Ethernet-like interface MIB
- RFC 2787: VRRP MIB
- RFC 2819: RMON MIB
- RFC 2863: Interface Group MIB
- RFC 2863: Interface MIB
- RFC 2922: LLDP MIB
- RFC 2925: Ping/Traceroute 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: SNMP Application 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
- RFC 4188: STP and Extensions MIB
- RFC 4363: Definitions of Managed Objects for Bridges with traffic classes, multicast filtering, and VLAN extensions
- Draft-ietf-idr-bgp4-mibv2-02.txt: Enhanced BGP-4 MIB
- Draft-ietf-isis-wg-mib-07
- Draft-reeder-snmpv3-usm-3desede-00
- Draft-ietf-idmr-igmp-mib-13
- Draft-ietf-idmr-pim-mib-09
- Draft-ietf-bfd-mib-02.txt

Network Management—MIB Support

- RFC 1155: Structure of Management Information (SMI)
- RFC 1157: SNMPv1
- RFC 1212, RFC 1213, RFC 1215: MIB-II, Ethernet-like MIB, and traps
- RFC 1657: BGP-4 MIB
- RFC 1724: RIPv2 MIB
- RFC 1850: OSPFv2 MIB
- RFC 1901: Introduction to Community-based SNMPv2
- RFC 1902: Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2)
- RFC 1905, RFC 1907: SNMP v2c, SMIv2, and Revised MIB-II
- RFC 2011: SNMPv2 for IP using SMIv2
- RFC 2012: SNMPv2 for transmission control protocol using SMIv2

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 10,000 ft (3,048 m)
- Relative humidity operating: 5 to 90% (noncondensing)
- Relative humidity nonoperating: 5 to 95% (noncondensing)
- Seismic: Designed to meet GR-63, Zone 4 earthquake requirements

Maximum Thermal Output

(estimated, subject to change)

- QFX10008: 42,148 BTU/hour (12,360 W)
- QFX10016: 88,316 BTU/hour (25,899 W)

Safety and Compliance

Safety

- CAN/CSA-C22.2 No. 60950-1 Information Technology Equipment—Safety
- UL 60950-1 Information Technology Equipment—Safety
- EN 60950-1 Information Technology Equipment—Safety
- IEC 60950-1 Information Technology Equipment—Safety (All country deviations)
- EN 60825-1 Safety of Laser Products—Part 1: Equipment Classification

Electromagnetic Compatibility

- 47CFR Part 15, (FCC) Class A
- ICES-003 Class A
- EN 55022 Class A
- CISPR 22 Class A
- EN 55024
- CISPR 24
- EN 300 386
- VCCI Class A
- AS/NZA CISPR22 Class A
- KN22 Class A
- 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



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. 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

Model Number	Description
QFX10000 Hardware	
QFX10008-REDUND	QFX10008 redundant 8-slot chassis with 2 CBs, 6 2,700 W AC power supplies, 2 fan trays, 2 fan tray controllers, and 6 switch fabric cards
QFX10008-BASE	QFX10008 Base 8-slot chassis with 1 CB, 3 2,700 W AC power supplies, 2 fan trays, 2 fan tray controllers, and 5 switch fabric cards
QFX10008-REDUND-DC	QFX10008 Redundant 8-slot chassis with 2 CBs, 6 2,500 W DC power supplies, 2 fan trays, 2 fan tray controllers, and 6 switch fabric cards
QFX10016-REDUND	QFX10016 Redundant 16-slot chassis with 2 CBs, 10 2,700 W AC power supplies, 2 fan trays, 2 fan tray controllers, and 6 switch fabric cards
QFX10016-BASE	QFX10016 Base 16-slot chassis with 1 CB, 5 2,700 W AC power supplies, 2 fan trays, 2 fan tray controllers, and 5 switch fabric cards
QFX10016-REDUND-DC	QFX10016 Redundant 16-slot chassis with 2 CBs, 10 2,500 W DC power supplies, 2 fan trays, 2 fan tray controllers, and 6 switch fabric cards

Model Number	Description
QFX10000 Line Cards	
QFX10000-36Q	QFX10000 36-port 40GbE QSFP+ / 12-port 100GbE QSFP28 line card
QFX10000-30C	QFX10000 30-port 100GbE QSFP28/40GbE QSFP+ line card
QFX10000-60S-6Q	QFX10000 60-port 1/10GbE SFP/SFP+ line card with 6 40GbE QSFP+ / 2 100GbE QSFP28 ports
QFX10000 Pluggable Optics	
JNP-QSFP-100G-SR4	QSFP28 100GBASE-SR4 Optics for up to 100 m transmission over parallel multi-mode fiber (MMF)
JNP-QSFP-100G-LR4	QSFP28 100GBASE-LR4 Optics for up to 10 km transmission over serial single-mode fiber (SMF)
JNP-QSFP-100G-PSM4*	QSFP28 100GBASE-PSM4 Optics for up to 2 km transmission over parallel SMF
JNP-QSFP-100G-CWDM*	QSFP28 100GBASE-CWDM4 Optics for up to 2 km transmission over serial SMF
JNP-QSFP28-AOC-10M*	100G QSFP28 to QSFP28 active optical cables, 10 m
JNP-QSFP-40GE-IR4	QSFP+ 40GBASE-IR4 40 Gigabit optics, 1,310 nm for up to 2 km transmission on SMF
JNP-QSFP-40G-LR4	QSFP+ 40GBASE-LR4 40 Gigabit optics, 1,310 nm for up to 10 km transmission on SMF
JNP-QSFP-40G-LX4	QSFP+ 40GBASE-LX4 40 Gigabit optics, 100 m (150 m) with OM3 (OM4) duplex MMF fiber
JNP-QSFP-4X10GE-IR	QSFP+ 40GBASE Optics, up to 1.4 km transmission on parallel single mode (4X10GbE long reach up to 1.4 km)
JNP-QSFP-4X10GE-LR	QSFP+ 40GBASE optics, up to 10 km transmission on parallel single mode (4X10GbE long reach up to 10 km)
QFX-QSFP-40G-ESR4	QSFP+ 40GBASE-ESR4 40 Gigabit optics, 300 m (400 m) with OM3 (OM4) MMF
QFX-QSFP-40G-SR4	QSFP+ 40GBASE-SR4 40 Gigabit optics, 850 nm for up to 150 m transmission on MMF
QFX-SFP-10GE-ER	SFP+ 10GBASE-ER 10-Gigabit Ethernet optics, 1,550 nm for 40 km transmission on SMF
QFX-SFP-10GE-LR	SFP+ 10GBASE-LR 10-Gigabit Ethernet optics, 1,310 nm for 10 km transmission on SMF
QFX-SFP-10GE-SR	SFP+ 10GBASE-SR 10-Gigabit Ethernet optics, 850 nm for up to 300 m transmission on MMF
QFX-SFP-10GE-USR	SFP+ 10-Gigabit Ethernet ultra short reach optics, 850 nm for 10 m on OM1, 2 0 m on OM2, 100 m on OM3 multi-mode fiber
QFX-SFP-1GE-LX	SFP 1000BASE-LX Gigabit-Ethernet optics, 1,310 nm for 10 km transmission on SMF
QFX-SFP-1GE-SX	SFP 1000BASE-SX Gigabit-Ethernet optics, 850 nm for up to 550 m transmission on MMF
QFX-SFP-1GE-T	SFP 1000BASE-T Copper transceiver module for up to 100 m transmission on Cat5

*Roadmap

Model Number	Description
QFX10000 Software Feature Licenses	
QFX10008-PFL	QFX10008 Premium Feature License
QFX10008-AFL	QFX10008 Advanced Feature License
QFX10016-PFL	QFX10016 Premium Feature License
QFX10016-AFL	QFX10016 Advanced Feature License
QFX10000 Power Cords	
CBL-EX-PWR-C19-AR	AC Power Cable - Argentina (16A/250V, 2.5m)
CBL-EX-PWR-C19-AU	AC Power Cable - Australia (15A/250V, 2.5m)
CBL-EX-PWR-C19-BR	AC Power Cable - Brazil (16A/250V, 2.5m)
CBL-EX-PWR-C19-C20	AC Patch Cable - C20 PDU (16A/250V, 2.5m)
CBL-EX-PWR-C19-CH	AC Power Cable - China (16A/250V, 2.5m)
CBL-EX-PWR-C19-EU	AC Power Cable - EU (16A/250V, 2.5m)
CBL-EX-PWR-C19-IL	AC Power Cable - Israel (16A/250V, 2.5m)
CBL-EX-PWR-C19-IN	AC Power Cable - India (16A/250V, 2.5m)
CBL-EX-PWR-C19-INT	AC Power Cable - Int'l/UK (IEC 013-6 16A/250V, 2.5m)
CBL-EX-PWR-C19-IT	AC Power Cable - Italy (16A/250V, 2.5m)
CBL-EX-PWR-C19-JP	AC Power Cable - Japan (16A/250V, 6-20, 2.5m)
CBL-EX-PWR-C19-JPL	AC Power Cable - Japan Locking (16A/250V, L6-20, 2.5m)
CBL-EX-PWR-C19-KR	AC Power Cable - Korea/EU (16A/250V, 2.5m)
CBL-EX-PWR-C19-SA	AC Power Cable - South Africa (16A/250V, 2.5m)
CBL-EX-PWR-C19-SZ	AC Power Cable - Switzerland (16A/250V, 2.5m)
CBL-EX-PWR-C19-TW	AC Power Cable - Taiwan (16A/250V, 2.5m)
CBL-EX-PWR-C19-UK	AC Power Cable - UK (13A/250V, 2.5m)
CBL-EX-PWR-C19-US	AC Power Cable - US (16A/250V, NEMA 6-20, 2.5m)
CBL-EX-PWR-C19-USL	AC Power Cable - US Locking (16A/250V, NEMA L6-20, 2.5m)
QFX10000 Field Replaceable Units	
QFX10008-CHAS	QFX10008 Chassis, spare
QFX10008-FAN	QFX10008 Fan, spare
QFX10008-SF	QFX10008 Switch fabric card, spare
QFX10008-FAN-CTRL	QFX10008 Fan tray controller, spare
QFX10016-CHAS	QFX10016 Chassis, spare
QFX10016-FAN	QFX10016 Fan, spare
QFX10016-SF	QFX10016 Switch fabric card, spare
QFX10016-FAN-CTRL	QFX10016 Fan tray controller, spare

Model Number	Description
QFX10000-RE	QFX10000 Control Board, spare
QFX10000-PWR-AC	QFX10000 2,700 W AC Power Supply, spare
QFX10000-PWR-DC	QFX10000 2,500 W DC Power Supply, spare
QFX10000-RE-BLNK	QFX10000 Control Board blank cover panel, spare
QFX10000-LC-BLNK	QFX10000 Line card blank cover panel, spare
QFX10000-PWR-BLNK	QFX10000 Power supply blank cover panel, spare

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.
 1133 Innovation Way
 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

Copyright 2016 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.

JUNIPER
 NETWORKS