


QFabric System



Product Overview

Modern data centers demand scale, performance, flexibility, and control—along with reductions in total cost of ownership (TCO). Juniper Networks QFabric architecture is a revolutionary approach that provides a foundation for cloud-ready virtualized data center network environments.

Juniper Networks QFabric System is the cornerstone of the QFabric architecture—a purpose-built solution that allows the creation of high-performance, scalable, cost-effective, dynamic, and easy to manage large enterprise IT or service provider cloud data centers.

Product Description

Juniper Networks® QFabric™ System is the only fabric solution that delivers any-to-any connectivity and simplified operations, making it the ideal architectural foundation for virtualized data centers today and for the next decade. It is a scalable, high-performance, non-blocking, and easy-to-manage fabric that enables traditional Layer 2 and Layer 3 connectivity along with virtualization and convergence. The standards-based QFabric System is completely interoperable and seamlessly integrates with customers' existing data center environments, allowing them to easily migrate traditional tiered networks to a single tier QFabric architecture that connects compute, storage, network, and services resources as extensions of a low latency network.

QFabric technology enables customers to maximize the performance of their data centers and simplify their network operations. By providing direct connectivity and predictable high performance at scale between any two ports in the fabric, common changes in the data center such as adding capacity, virtual machine mobility, or deploying new applications can be achieved quickly and easily.

Two QFabric System models are available:

- The QFX3000-M QFabric System, designed for mid-tier, satellite, and container data center environments, supports from 48 to 768 10GbE ports, delivering the simplicity, agility, and performance benefits of QFabric architecture in a space optimized form factor. The QFX3000-M is ideal for High Performance Computing environments, Big Data Hadoop clusters, and back-end business applications, and it provides investment protection by allowing customers to easily scale to a larger QFabric system deployment as demands for 10GbE grow.
- The QFX3000-G QFabric System, designed for large enterprises, service providers, and cloud data center environments, scales to support up to 6,144 10GbE ports in a single device. The QFX3000-G is ideal for cloud (IaaS, SaaS), large enterprise IT data center (business applications, data analytics), and High Performance Computing (grid computing, data modeling, scientific research) environments.

QFabric System Components

The QFabric System consists of three separate but interdependent edge, interconnect, and control devices—the QFabric Node, QFabric Interconnect, and QFabric Director. These represent the internal elements of a traditional switch.

- **QFabric Node:** In a QFabric system, the line cards that typically reside within a modular chassis switch become high-density, fixed-configuration, 1 U edge devices that provide access into and out of the fabric. The Nodes, which can also operate as independent top-of-rack 10GbE switches*, provide compute, storage, services, and network access for the QFabric System. There are two types of QFabric Nodes available: the QFX3500, which offers a variety of connectivity options ranging from 1GbE to 10GbE, Fibre Channel (FC), and FC over Ethernet (FCoE); and the QFX3600, which offers 10GbE and 40GbE* connectivity options. Both the QFX3500 and QFX3600 Nodes can be used in a single system.

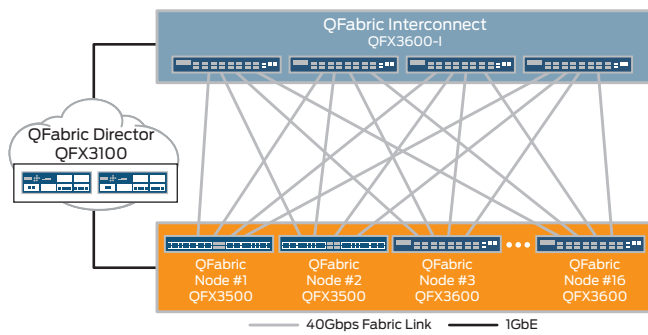


Figure 1: The QFX3000-M system topological view

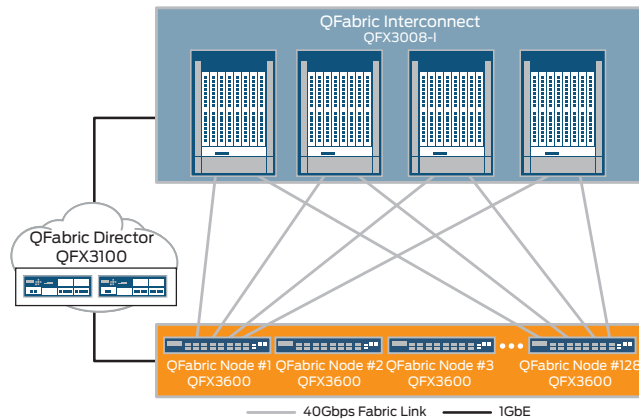


Figure 2: The QFX3000-G system topological view

- **QFabric Interconnect:** The QFabric Interconnect represents the typical backplane of a modular switch, connecting all QFabric Node edge devices in a flat, any-to-any topology. This topology provides the data plane connectivity between all Nodes, with the Interconnect acting as the high-performance backplane. Two QFabric Interconnect options are available. The QFX3000-M uses the 1 U fixed configuration QFX3600-I QFabric Interconnect, which supports up to 16 connected QFabric Nodes to create a single fabric capable of supporting 768 10GbE ports. The QFX3000-G uses the modular QFX3008-I, which connects up to 128 QFabric Nodes to create a single fabric capable of supporting 6,144 10GbE ports.
- **QFabric Director:** The Routing Engines embedded within a modular switch are externalized in the QFabric system via the QFX3100 QFabric Director, which provides control and management services for the fabric. Deployed in clusters to provide redundancy, QFabric Directors provide a single management interface to manage the scalable data plane provided by the Node and Interconnect devices.

The QFabric Node and QFabric Interconnect devices together create the distributed data plane for the QFabric System over which all data traffic to and from servers and storage is carried. Existing QFabric system components can be redeployed between a QFX3000-M and a QFX3000-G, greatly simplifying flexibility and migration. Users can initially deploy a QFX3000-M and, as their 10GbE demands grow, migrate to a QFX3000-G with the simple replacement of the QFabric Interconnect, dramatically increasing scale.

* Roadmap for QFX3600

One of the greatest advantages of QFabric technology is its manageability. Unlike traditional deployments with multiple touch points for provisioning and troubleshooting, a QFabric System presents a single management interface for provisioning, managing, and troubleshooting the data center. Up to 128 top-of-rack switches in a QFX3000-G system and up to 16 top-of-rack switches in a QFX3000-M system work together to connect network, compute, and storage resources.

Architecture and Key Components

With the Node, Interconnect, and Director, the QFabric System operates as a single logical device—a distributed system with a separate data and control plane. These independent devices are all managed from a single point, greatly simplifying the provisioning, monitoring, and troubleshooting of the network fabric.

The QFabric architecture, as implemented in a QFabric System, is composed of three distinct planes—the data plane, the control plane, and the management plane.

- **Data plane:** The QFabric Node and QFabric Interconnect devices create the distributed data plane for the QFabric architecture. All data traffic to and from servers and storage is carried over the data plane. Connectivity between Node and Interconnect devices is provided via two, four, or eight 40 Gbps links. All links provide active connectivity between end points, eliminating the need for Spanning Tree Protocol (STP). L2, L3, and FCoE traffic is effectively load-balanced across all available links between Nodes and Interconnects. Single or multiple links from Nodes to each Interconnect can be used, and up to four Interconnects can be used in a QFX3000-M or QFX3000-G system.
- **Control plane:** Separating the data plane from the control plane has always been a key design principle for Juniper to ensure high reliability. This principle is incorporated in the QFabric architecture, where data and control traffic are carried over two separate networks. The distributed nature of the control plane is the key to the QFabric architecture's scalability and simplicity, providing tremendous reliability by eliminating any single point of failure in the system. Control plane services are provided by the QFabric Director, which uses a completely separate out-of-band control plane network that connects Directors and Nodes to the Interconnect and is used exclusively for carrying control traffic. The Juniper Networks EX4200 Ethernet Switch is typically used to connect QFabric Directors to QFabric Interconnects and QFabric Nodes over a 1GbE out-of-band control plane network.
- This out-of-band network is used for auto-discovery of all devices, provisioning, image upgrades for various elements in the system, and configuration. All of these functions are fully automated and do not require user configuration.
- **Management plane:** The Director provides all management services for the QFabric architecture, communicating directly with all Node and Interconnect devices to build

a global view of the entire network. This provides a single point of visibility, control, and management for the entire data center fabric, and it significantly reduces the operational costs typically associated with managing a data center network. The Director also interfaces with the network management ecosystem via standards-based protocols such as XML/NETCONF, SNMP, or command-line interface (CLI). Junos XML management protocol and Juniper Networks Junos® SDK provide a rich automation framework that enables network customization and tuning as required, ensuring that the QFabric architecture deployment fits the existing ecosystem without having to invest in special tools.

High-Performance Layer 2/Layer 3 Deployments

The QFabric System is designed to provide a low latency fabric that can scale to more than 6,000 ports and be deployed in a variety of environments. With the advent of server virtualization, the IT infrastructure is providing business efficiency by consolidating many physical servers into fewer high-performance virtualized servers. However, this introduces new challenges in the data center by significantly increasing network utilization and requiring faster access-layer connectivity. Every QFabric Node in a QFabric System adds high-performance, ultra-low latency (ULL) 10GbE ports, making it possible to support large-scale server virtualization deployments with a large media access control (MAC) address table with ultra-low latency (5 microseconds port-to-port under typical loads for a QFX3000-G system, and 3 microseconds port-to-port under typical loads for a QFX3000-M system) at L2 and L3 from server node to server node.

The QFabric system offers the following advantages for high-performance access:

- Full featured, standards-based L2 and L3 switching capabilities
- Low latency switching on up to 56 10GbE ports with the QFX3600 Node or 48 10GbE ports with the QFX3500 Node
- Scaling options for 768 10GbE ports with the QFX3000-M system or 6,144 10GbE ports with the QFX3000-G system using QFX3500 or QFX3600 Nodes at 3:1 or 6:1 oversubscription
- Scaling options for up to 896 10GbE ports with the QFX3000-M system or 7,168 10GbE ports with the QFX3000-G system using QFX3600 nodes at 7:1 oversubscription
- Support for the same Juniper Networks Junos® operating system that powers other Juniper Networks switches, routers, and security products, as well as the Juniper Networks Junos® Space management platform

Virtualization and I/O Convergence

As businesses adopt 10GbE connectivity in the data center access layer, CapEx reduction can be achieved by consolidating storage and Ethernet traffic on common 10GbE server connections. The

deployment of virtualized servers hosting hundreds of virtual machines and high-performance servers with converged network adapters, as well as storage and I/O convergence, will require low latency and lossless 10GbE technologies at the server access layer to support FC and FCoE interfaces.

For end-to-end convergence, the QFabric architecture offers extensive Data Center Bridging (DCB) capabilities, including specific iSCSI support. For FCoE-based converged server edge access environments, the QFabric System can also operate as an FCoE transit switch and FCoE-to-Fibre Channel (FCoE-FC) gateway, enabling customers to protect their investments in existing data center aggregation and FC storage area network (SAN) infrastructures.

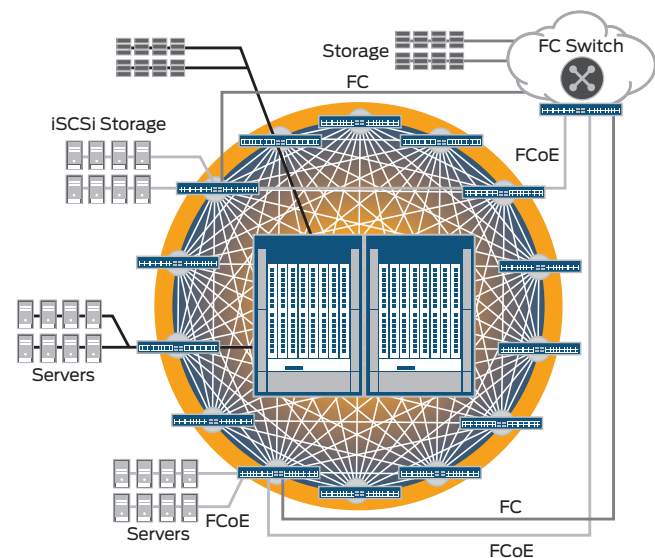


Figure 3: The QFX3000-G as an FCoE iSCSI transit switch or FCoE-to-FC gateway

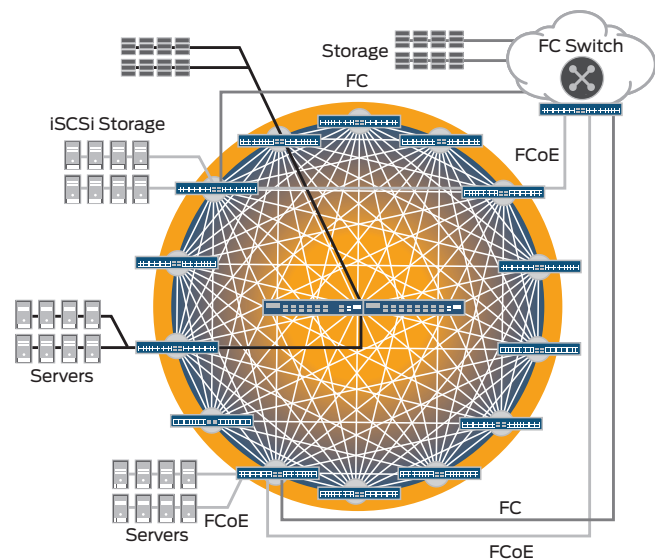


Figure 4: The QFX3000-M as an FCoE iSCSI transit switch or FCoE-to-FC gateway

Cloud-Ready Data Center Network

Cloud architectures—both private and public—require high levels of scalability, elasticity, and multi-tenancy capabilities. For many other enterprises, operational efficiency is absolutely critical.

The QFabric System can scale from just a few hundred ports to thousands of server/storage ports, helping customers build highly scalable, high performing, highly efficient, and cloud-ready (private, public, or hybrid) data center infrastructures.

QFabric System Features and Benefits

Feature	Benefit
Multi-terabit capacity and scalability	The QFabric architecture allows for scale and incremental growth by adding a QFabric Node when a new server rack is installed. With the QFabric Interconnect acting as the network backplane with a capacity of 5.12 Tbps in a QFX3000-M system and 40 Tbps in a QFX3000-G system, the QFabric architecture scales to support up to 6,144 10GbE ports, with the ability to scale beyond 10GbE in the future for both types of systems. The incremental investment required to connect new server and storage capacity is a fraction of the time and cost of the existing tiered architecture, and returns on investments are much faster than existing deployment models.
Low latency fabric	The QFabric architecture's any-to-any connectivity supports high-speed server-to-server communication. QFabric technology has interface-to-interface latency on the order of 900 nanoseconds, and 3 to 5 microseconds under typical loads across the fabric between any two interfaces for the QFX3000-M and QFX3000-G, respectively. Ultra-low latency provides an order of magnitude improvement in performance over traditional network architectures, making the QFabric architecture ideal for supporting latency sensitive applications, east-west traffic flows, virtualization, cloud computing, and other high-performance data center initiatives.
Virtualization and converged I/O architecture	The QFabric architecture is optimized for server virtualization, making it easier for enterprises to eliminate "stranded capital" resulting from the underutilization of existing server and storage assets. The QFabric architecture also supports converged traffic at 10GbE access port speeds, enabling enterprises to migrate to 10GbE converged access on a single network. QFabric technology eliminates the need to provision separate networks for LAN and SAN trade-offs that three-tiered tree networks must make with respect to simplicity, performance, scalability, and cost while delivering the same scale, performance, and enhanced features as traditional networks.
Robust, modular, feature-rich software	Each release of Junos OS runs consistently across all Juniper Networks routing platforms and feature sets. Junos OS was conceived and implemented as a modular design, with each Junos OS process running in protected memory to guard against system crashes and to ensure that applications do not interfere with each other. Junos OS provides the greatest breadth of features and most stable network operating system in the industry.
Operational performance	By presenting itself as a single switch running Junos OS, the QFabric system greatly simplifies data center management. And by reducing the number of switches in the data center network, the QFabric architecture significantly lowers complexity and operational expenses, as well as power, space, and cooling costs. In addition, to maximize network uptime, Juniper has designed the QFabric architecture as a reliable carrier-class infrastructure with no single point of failure or downtime for reconfiguration and maintenance.
Interface flexibility	The QFabric System provides a range of interfaces for next-generation data center access from 1GbE, 10GbE, and 40GbE* for Ethernet. The QFabric architecture also provides FCoE as well as the FC interfaces required for converged I/O environments. This interface variety (both optical and copper) in a single tier network offers distinct advantages in terms of OpEx and CapEx reduction. Combining the functions of previously disparate compute, storage, and services offers greater network simplicity and retains the service building advantages of the converged network.
Scaling efficiency/increased ROI	QFabric architecture provides for significant investment protection and enhances the "pay-as-you-grow" model by eliminating the need for network tiers. Adding compute clusters simply requires deploying QFX3500 or QFX3600 Nodes to the fabric. This does not require upfront investment and growth can be spread over multiple phases. As new points of delivery are added to the data center, QFabric Nodes can be connected and provisioned into the existing QFabric architecture, reducing CapEx and OpEx while providing efficiencies of scale and performance.
Green	QFabric architecture is environmentally conscious, allowing enterprises to optimize every facet of the data center network while consuming less power, requiring less cooling, and producing a fraction of the carbon footprint of legacy multitiered data center networks. The individual components of QFabric technology are also designed to meet green standards such as Restriction of Hazardous Substances (RoHS), Registration, Evaluation, Authorisation and Restriction of Chemical substances (REACH), 80 Plus, Waste Electrical and Electronic Directive (WEEE Directive), and others.

* Roadmap

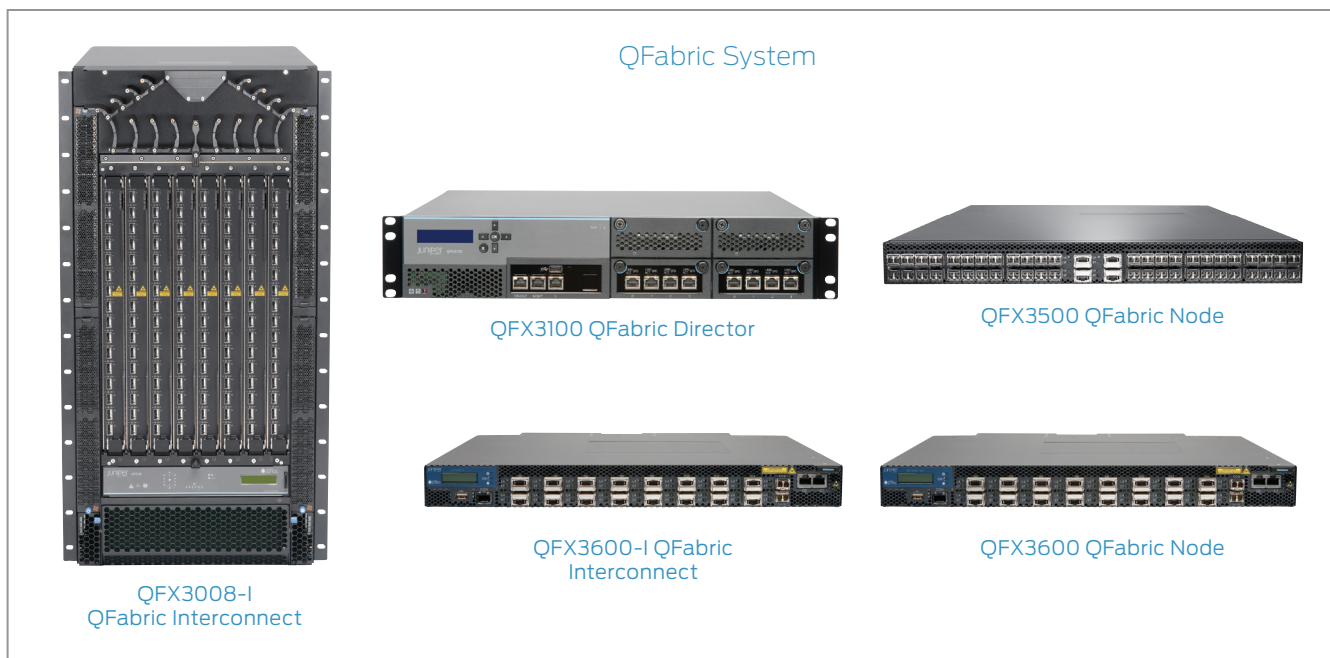
Junos Operating System

QFabric system components run the same reliable and high-performance Junos OS that is used by Juniper Networks EX Series Ethernet Switches, Juniper routers, and Juniper Networks SRX Series Services Gateways. By utilizing a common operating system, Juniper delivers a consistent implementation and operation of control plane features across products. To maintain that consistency, Junos OS adheres to a highly disciplined development process, follows a single release track, and employs a highly available modular architecture that prevents isolated failures from bringing down an entire system. These attributes are fundamental to the core value of the software, enabling Junos OS-powered products to be updated simultaneously with the same software release. Features are fully regression tested, making each new release a superset of the previous version. Customers can deploy the software with confidence that existing capabilities will be maintained and operate in the same way.

QFabric Architecture Management Capabilities

The following system management options are available for the QFabric system:

- The standard Junos OS CLI and SNMP module offers the same granular management capabilities and scripting parameters found in any router or switch powered by Junos OS.
- Performance, configuration, and fault data for the QFabric architecture 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.
- QFabric technology is supported by Junos Space, an open, programmable application platform for hosting a comprehensive suite of network operational application tools that provide a smart, simple, and open approach for automating the deployment and operation of a Juniper infrastructure.
- QFabric architecture also supports Junos XML management protocol, which facilitates application and script development to develop custom applications easily and quickly using the XML programming language. Junos XML management protocol automation tools also provide early detection and automatic resolution of potential problems related to the operating system.



QFabric System Specifications

Attribute	QFX3500 QFabric Node	QFX3600 QFabric Node	QFX3600-I QFabric Interconnect	QFX3008-I QFabric Interconnect	QFX3100 QFabric Director
Width	17.5 in. (44.45 cm)	17 in. (43.2 cm)	17 in. (43.2 cm)	17.5 in. (44.45 cm)	17.5 in. (44.45 cm)
Height	1.75 in. (4.45 cm), 1 U	1.74 in. (4.4 cm) 1 U	1.74 in. (4.4 cm) 1 U	36.75 in. (93.34 cm), 21 U	3.5 in. (8.89 cm), 2 U
Depth	28 in. (71.12 cm)	19.4 in. (49.3 cm)	19.4 in. (49.3 cm)	32 in. (81.28 cm)	23.75 in. (60.33 cm)
Weight	30.75 lb (13.95 kg)	30.8 lb (14 kg)	30.8 lb (14 kg)	675 lb (306.17 kg) fully populated	41.2 lb (18.69 kg)
Power feed (voltage)	100–240 V AC (single phase) –40 to –72 V DC	100–240 V AC (single phase) –40 to –72 V DC	100–240 V AC (single phase) –40 to –72 V DC	200–240 V AC (single phase); 240 V (three phase)	100–240 V AC (single phase)
Power feed (AMP rating)	7.8 A (100–127 V) 3.8 A (200–240 V)	7.8 A (100–127 V) 3.8 A (200–240 V)	7.8 A (100–127 V) 3.8 A (200–240 V)	Single phase: 16 A per input Three phase: 25.5 A per phase	8 A
Power consumption (maximum)	365 watts	345 watts	345 watts	6,240 watts (fully loaded)	476 watts
Power consumption (nominal)	230 watts	255 watts	255 watts	4,620 watts (fully loaded)	220 watts
Heat dissipation (maximum)	1,250 BTU	1,177 BTU	1,177 BTU	21,290 BTU	1,624 BTU
Heat dissipation (nominal)	784 BTU	870 BTU	870 BTU	15,763 BTU	751 BTU
Air flow direction	Front to back Back to front	Front to back Back to front	Front to back Back to front	Front to back	Front to back
Rack mount options	4–pole rack mount	4–pole rack mount 2–pole mid mount	4–pole rack mount 2–pole mid-mount	4–pole rack mount	4–pole rack mount 2–pole mid mount
Operating altitude range	13,000 ft (3,962 m)	13,000 ft (3,962 m)	10,000 ft (3,048 m)	13,000 ft (3,962 m)	10,000 ft (3,048 m)

Layer 2 Features

- VLAN—IEEE 802.1Q VLAN trunking
- Routed VLAN interface (RVI)
- Port-based VLAN
- MAC address filtering
- Static MAC address assignment for interface
- Per VLAN MAC learning (limit)
- Link aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
- QinQ Tunneling
- MVRP
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
- Jumbo frame (9,216 bytes)

Link Aggregation

- LAG load-sharing algorithm—bridged or routed (unicast or multicast) traffic
 - IP: Source IP (SIP), Destination IP (DIP), TCP/UDP source port, TCP/UDP destination port
 - Layer 2 and non-IP: MAC source address, MAC destination address, Ethertype, VLAN ID, source port
 - FCoE packet: service identifier (SID), direct inward dialing (DID), OxID, source port

Layer 3 Features (IPv4)

- Static routing
- Routing policy
- Routing protocols (OSPF, BGP, IS-IS*, RIP*, etc.)
- Routed ports, RVIs, and L3 LAG
- Virtual routing and forwarding (VRF): VRF-lite, VRF-aware unicast (BGP, OSPF)
- Virtual Router Redundance Protocols (VRRP)

Multicast Features

- Internet Group Management Protocol (IGMP) snooping v1 and v2
- IGMP snooping v3*
- Layer 3 multicast routing protocols
- Multicast Source Directory Protocol (MSDP)
- Protocol Independent Multicast Version 2 (PIMv2) Sparse Mode (PIM-SM) and Source-Specific multicast (SSM)*
- Bootstrap router (BSR), auto-RP, and static RP

Security and Firewall Filters (ACLs)

- Secure interface login and password
- RADIUS
- TACACS+
- Ingress and egress firewall filters—allow and deny, port ACLs, VLAN ACLs, routed ACLs

- ACL actions—logging, system logging, reject, mirror to an interface, counters, assign forwarding class, permit, drop, police, mark
- SSH v1, v2
- Local proxy Address Resolution Protocol (ARP)
- Static ARP support
- Storm control, port error disable, and auto-recovery
- Control plane denial-of-service (DoS) protection

Quality of Service (QoS)

- Layer 2 QoS—classification, rewrite, queuing
- Layer 3 QoS
- Rate limiting
 - Ingress/egress policing—1 rate 2 color, 1 rate 3 color, 2 rate 3 color
 - Egress policer—policer mark down action
 - Egress shaping—per queue, per port
- Twelve hardware queues per port per node (8 unicast and 4 multicast)
- Strict-priority queue (SPQ), shaped-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/Dynamic Host Configuration Protocol (DSCP) (ingress)
- Remarking of bridged packets

Traffic Mirroring

- Port-based
- LAG port
- VLAN-based
- Filter-based
- Mirror to local and remote destinations (L2 over VLAN)
- sFlow technology

Data Center Bridging (DCB)

- Priority Flow Control (PFC)—IEEE 802.1Qbb
- Enhanced Transmission Selection (ETS)—IEEE 802.1Qaz
- Ethernet Congestion Management (QCN)—IEEE 802.1Qau*
- 802.1Qbg VEPA Hairpin Switching
- Data Center Bridging Exchange Protocol (DCBX)—part of the ETS standard
 - 1.01 mode
 - IEEE mode
 - FCoE application TLV
 - iSCSI application TLV
 - User configurable application TLV

Fibre Channel over Ethernet (FCoE)

- FCoE transit switch (FIP snooping)
- FCoE-FC gateway
- iSCSI transit switch (iSCSI tlV)

Fiber Channel Standard

- Fibre Channel port speeds—2, 4, 8 Gbps
- Fibre Channel port types—N_Port and VF_Port (fabric only mode)
- Fibre Channel classes of service—Class 3
- Fibre Channel services—N_Port Virtualizer Device (FCoE to FC)
- Fibre Channel services— N_Port ID Virtualization (NPIV) gateway
- FCoE Support—FC-BB-5 FC-BB_E including FCoE Initialization Protocol (FIP) Snooping

QFabric Architecture Server Virtualization Management

- Junos Space Network Director
- IEEE 802.1Qbg*

QFabric Architecture Management and Operations

- Fabric visualization with Network Director
- Role-based CLI management and access
- CLI via console and SSH
- Show and debug commands, statistics
- Fabric ping and traceroute
- Junos OS configuration rescue and rollback
- SNMP v1/v2 and v3*
- XML/NETCONF
- Junos XML management protocol automation
- Junos SDK
- Rolling software upgrade

QFabric System Performance and Scale

QFabric system performance and scale can be characterized by the cumulative resources of the Nodes being used for the network design. Following are the scale numbers for the QFabric architecture:

Platform:

- QFabric Directors: 2
- QFabric Interconnects: 4 QFX3600-I in QFX3000-M system;
4 QFX3008-I in QFX3000-G system
- QFabric Nodes: 16 QFabric Nodes in QFX3000-M system;
128 QFabric Nodes in QFX3000-G system
- QFX3000-M port scale options
 - 768 10GbE ports with QFX3500 or QFX3600 Nodes at 3:1 or 6:1 oversubscription
 - 896 10GbE ports in QFX3000-M system with QFX3600 Nodes at 7:1 oversubscription
- QFX3000-G port scale options
 - 6,144 10GbE ports with QFX3500 or QFX3600 Nodes at 3:1 or 6:1 oversubscription
 - 7,168 10GbE ports in QFX3000-G system with QFX3600 Nodes at 7:1 oversubscription

* Roadmap

Layer 2

- MAC addresses: 128,000 through 1,536,000
- VLANs: 4,096 active; 4,090 configurable
- Number of LAGs: 48 per QFabric Node
- Number of ports per LAG: 32
- Jumbo frame: 9,216 bytes

Layer 3

- RVIs: 2,000
- IPv4 unicast routes: 16,000 for QFabric architecture
- OSPF neighbors: 256
- BGP peers: 256

Multicast

- Multicast groups: 4,000 for QFabric architecture

QoS

- Policers (ingress and egress): 1,500 per QFabric Node
- Queues: 12 per QFabric Node, 8 unicast and 4 multicast

Security

- Firewall filters (ACLs): 1,500 per QFabric Node
- Traffic mirroring
 - Mirroring destination ports per QFabric Node: 4
 - Mirroring destination VLANs per QFabric Node: 256

FCoE/FC

- FCoE interfaces: 6,144 maximum
- FC interfaces: 1,536 maximum
- FC sessions: 3,000 per QFabric Node
- VF_Ports: 6,144 maximum

QFabric Architecture Standards Compliance

IEEE Standard

- IEEE 802.1AB: Link Layer Discovery Protocol (LLDP)
- IEEE 802.1p: Class-of-service (CoS) prioritization
- IEEE 802.1Q-2006: VLAN tagging
- IEEE 802.3ab: 1000BASE-T
- IEEE 802.3z: 1000BASE-X
- IEEE 802.3ae: 10-Gigabit Ethernet
- IEEE 802.3ad: Link Aggregation Control Protocol (LACP)
- IEEE 802.1Qbb
- IEEE 802.1Qaz
- IEEE 802.1Qau*
- IEEE 802.1Qbg*

T11 Standards:

- INCITS T11 FC-BB-5

Supported RFC

- RFC 768 UDP
- RFC 791 IP
- RFC 792 ICMP
- RFC 793 TCP
- RFC 826 ARP

- RFC 894 IP over Ethernet
- RFC 903 RARP
- RFC 1027 Proxy ARP
- RFC 1058 RIP v1
- RFC 1112 IGMP v1
- RFC 1122 Host requirements
- RFC 1195 Use of Open Systems Interconnection (OSI) IS-IS for Routing in TCP/IP and Dual Environments (TCP/IP transport only)
- RFC 1256 IPv4 ICMP Router Discovery (IRDP)
- RFC 1492 TACACS+
- RFC 1519 Classless Interdomain Routing (CIDR)
- RFC 1587 OSPF not-so-stubby area (NSSA) option
- RFC 1591 Domain Name System (DNS)
- RFC 1745 BGP4/IDRP for IP-OSPF interaction
- RFC 1765 OSPF database overflow
- RFC 1771 Border Gateway Protocol 4
- RFC 1812 Requirements for IP Version 4 routers
- RFC 1965 Autonomous system confederations for BGP
- RFC 1997 BGP communities attribute
- RFC 2030 SNTP, Simple Network Time Protocol
- RFC 2138 RADIUS Authentication
- RFC 2139 RADIUS accounting
- RFC 2154 OSPF with digital signatures (password, Message Digest 5)
- RFC 2236 IGMP v2
- RFC 2267 Network ingress filtering
- RFC 2328 OSPF v2 (edge mode)
- RFC 2362 PIM-SM (edge mode)
- RFC 2370 OSPF opaque link-state advertisement (LSA) option
- RFC 2385 TCP MD5 authentication for BGPv4
- RFC 2439 BGP Route flap damping
- RFC 2453 RIP v2
- RFC 2474 DiffServ precedence, including 8 queues/port
- RFC 2475 DiffServ core and edge router functions
- RFC 2597 DiffServ assured forwarding (AF)
- RFC 2598 DiffServ expedited forwarding (EF)
- RFC 2796 BGP route reflection (supersedes RFC 1966)
- RFC 2918 Route refresh capability for BGP-4
- RFC 3376 IGMP v3
- RFC 3392 Capabilities advertisement with BGP-4
- RFC 3569 Draft-ietf-ssm-arch-06.txt PIM-SSM PIM source-specific multicast
- RFC 3623 OSPF graceful restart
- RFC 4360 BGP extended communities attribute
- RFC 4486: Subcodes for BGP cease notification message
- Draft-ietf-idr-restart-10.txt: Graceful restart mechanism for BGP
- Draft-ietf-isis-restart-02: Restart signaling for IS-IS
- PIM-DM Draft IETF PIM: Dense mode draft-ietf-idmr-pimdm-05.txt, draft-ietf-pim-dm-new-v2-04.txt

* Roadmap

Supported MIBs

- RFC 1155 Structure and identification of management information for TCP/IP-based Internets
- RFC 1157 A Simple Network Management Protocol (SNMP)
- RFC1212 Concise MIB definitions
- RFC 1213 Management Information Base for network management of TCP/IP-based Internets: MIB-II (partial)
- RFC 1215 A convention for defining traps for use with the SNMP
- RFC 1901 Introduction to community-based SNMPv2
- RFC 1905 Protocol operations for Version 2 of the Simple Network Management Protocol (SNMPv2)
- RFC 1907 Management Information Base for Version 2 of the Simple Network Management Protocol (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 2571 An architecture for describing SNMP management frameworks (read-only access) (SNMPv3*)
- RFC 2572 Message processing and dispatching for the SNMP (read-only access) (SNMPv3*)
- RFC 2576 Coexistence between SNMP Version 1, Version 2, and Version 3 (SNMPv3*)
- RFC 2578 SNMP Structure of Management Information MIB
- RFC 2579 SNMP textual conventions for SMIv2
- RFC 2580 Conformance statements for SMIv2
- RFC 2863 Interface Group MIB
- RFC 3410 Introduction and applicability statements for Internet Standard Management Framework (SNMPv3*)
- RFC 3411 An architecture for describing SNMP management frameworks (SNMPv3*)
- RFC 3412 Message Processing and Dispatching for the SNMP (SNMPv3*)
- RFC 3413 Simple Network Management Protocol (SNMP) (all MIBs are supported except the Proxy MIB) (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 MIB for the SNMP (SNMPv3*)
- RFC 3584 Coexistence between Version 1, Version 2, and Version 3* of the Internet Standard Network Management Framework
- Support for the following Juniper Networks enterprise-specific MIBs:
 - Chassis MIB (jnx-chassis.mib)
 - Interface MIB (jnx-if-extensions.mib*)
 - Power supply unit MIB (jnx-power-supply-unit.mib)

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.

Model Number	Description
QFX3500 Node Base Hardware	
QFX3500-48S4Q-ACR	QFX3500, 48 small form-factor pluggable transceiver (SFP+/SFP) and 4 QSFP ports, redundant dual AC power supply, front-to-back air flow
QFX3500-48S4Q-AFI	QFX3500, 48 SFP+/SFP and 4 QSFP ports, redundant fan trays, field-replaceable unit (FRU) side to port side air flow (Note: Management module and power supplies are extra.)
QFX3500-48S4Q-AFO	QFX3500, 48 SFP+/SFP and 4 QSFP ports, redundant fan trays, port side to FRU side air flow (Note: Management module and power supplies are extra.)
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 Node Hardware Spare	
QFX3500-48S4Q	QFX3500, 48 SFP+/SFP and 4 QSFP ports (spare)
QFX3500-MB	Management board for QFX3500-48S4Q and QFX3500-48S4Q-AFI (spare)
QFX3500-FANAI	Fan module (intake) for QFX3500-48S4Q (spare)
QFX3500-RMAR	Rear rack mount bracket A (spare)
QFX3500-FAN-AFO	Fan module (port side to FRU side air flow) for QFX3500-48S4Q-AFO (spare)
QFX3500-MB-RJ45-AFO	Management board with 1GbE RJ45 interface for QFX3500-48S4Q-AFO (port side to FRU side air flow) (spare)
QFX3500-MB-SFP-AFI	Management board with SFP interface for QFX3500-48S4Q and QFX3500-48S4Q-AFI (FRU side to port side air flow) (spare)
QFX3500-MB-SFP-AFO	Management board with SFP interface for QFX3500-48S4Q-AFO (port side to FRU side air flow) (spare)
JPSU-650W-AC-AFI	650 W AC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow intake)
JPSU-650W-AC-AFO	650 W AC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow exhaust)

Model Number	Description
JPSU-650W-DC-AFI	650 W DC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow intake)
JPSU-650W-DC-AFO	650 W DC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow exhaust)

QFX3600 Node Base Hardware

QFX3600-16Q-AFI	QFX3600—16-port QSFP+ switch with three fans (FRU side to port side air flow); power supplies (2 required) and power cables (2 required) not included
QFX3600-16Q-AFO	QFX3600—16-port QSFP switch with three fans (port side to FRU side air flow); power supplies (2 required) and power cables (2 required) not included
QFX3600-16Q-ACR	QFX3600 16-port QSFP+ switch with three fans, FRU side to port side air flow and redundant AC power supplies
QFX3600-16Q-ACRB	QFX3600 16-port QSFP+ switch with three fans, port side to FRU side air flow and redundant AC power supplies

QFX3600 Node Hardware Spares

JPSU-650W-AC-AFI	650 W AC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow intake)
JPSU-650W-AC-AFO	650 W AC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow exhaust)
JPSU-650W-DC-AFI	650 W DC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow intake)
JPSU-650W-DC-AFO	650 W DC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow exhaust)

QFX3008-I Interconnect Base Hardware

QFX3008-CHASA-BASE	QFX3008-I, 8 slots AC base system with redundant dual control card, six redundant power supplies
QFX3008-SFI6Q	Front card with 16-port QSFP

QFX3008-I Interconnect Hardware Spares

QFXC08-ACTRAY-C19	A/C power wiring tray (single phase)
QFXC08-ACTRAY-D	A/C power wiring tray (three phase delta)
QFXC08-ACTRAY-W	A/C power wiring tray (three phase WYE)
QFXC08-CABMAN	Cable management module
QFXC08-DOOR	Front door
QFXC08-CHASA	QFX3008 chassis, 8 slots A/C
QFXC08-CB4S	Control card with 4x1/10GbE SFP+ ports
QFXC08-PWRAC-4000	4,000 W A/C PSU (spare)
QFXC08-FANT	Top fan tray (spare)
QFXC08-FANB	Bottom fan tray (spare)
QFXC08-FANS	Side fan tray (spare)
QFXC08-FBLNK	Front slot blank cover
QFXC08-FFLTR	Front air filter (spare)
QFXC08-SFLTR	Side air filter (spare)
QFX3008-SR1	Rear fabric card (spare)

Model Number	Description
--------------	-------------

QFX3600-I Interconnect Base Hardware

QFX3600-I-16Q-AFI	QFX3600-I QFabric Interconnect with three fans (FRU side to port side air flow); power supplies (two required) and power cables (two required) not included
QFX3600-I-16Q-AFO	QFX3600-I Interconnect with three fans (port side to FRU side air flow); power supplies (two required) and power cables (two required) not included
QFX3600-I-16Q-ACR	QFX3600-I Interconnect with three fans (FRU side to port side air flow) and redundant AC power supplies
QFX3600-I-16Q-ACRB	QFX3600-I Interconnect with three fans (port side to FRU side air flow) and redundant AC power supplies

QFX3600-I Interconnect Hardware Spares

JPSU-650W-AC-AFI	650 W AC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow intake)
JPSU-650W-AC-AFO	650 W AC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow exhaust)
JPSU-650W-DC-AFI	650 W DC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow intake)
JPSU-650W-DC-AFO	650 W DC power supply for EX4550, QFX3500, and QFX3600 (PSU side airflow exhaust)

QFX3100 Director Base Hardware

QFX3100-GBE-ACR	QFX3100 base system with redundant AC power supply, dual disks, and network interface cards
QFX3100-GBE-SFP-ACR	QFX3100 base system with redundant AC power supply, dual disks, and SFP network interface cards

QFX3100 Director Hardware Spares

QFX3100-NM-4GE	Network interface card with 4x10/100/1000BASE-TX (spare)
QFX3100-NM-4GE-SFP	Network interface card with 4x1GbE SFP (spare)
QFX3100-PWRAC-560A	560 W power supply A (spare)
QFX3100-HDD-2TB	2 TB hard drive (spare)
QFX3100-FANA	Fan module (spare)

Optics and Transceivers

QFX-QSFP-40G-SR4	QSFP+ 40GBASE-SR4 40-Gigabit Optics, 850 nm for up to 150 m transmission on multimode fiber-optic (MMF)
QFX-QSFP-40G-ESR4	QSFP+ 40GBASE-ESR4 40-Gigabit Optics, 300m (400m) with OM3 (OM4) MMF
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, 20 m on OM2, 100 m on OM3 MMF
QFX-SFP-10GE-LR	SFP+ 10GBASE-LR 10-Gigabit Ethernet Optics, 1,310 nm for 10 km transmission on single-mode fiber (SMF)

Model Number	Description
QFX-SFP-10GE-ER	SFP+ 10GBASE-ER 10-Gigabit Ethernet Optics, 1,550 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
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-SFP-1GE-T	SFP 1000BASE-T 10/100/1000 Copper Transceiver Module for up to 100 m transmission on Cat5
QFX-SFP-1GE-SX	SFP 1000BASE-SX Gigabit Ethernet Optics, 850 nm for up to 550 m transmission on MMF
QFX-SFP-1GE-LX	SFP 1000BASE-LX Gigabit Ethernet Optics, 1,310 nm for 10 km transmission on SMF
QFX-QSFP-DACBO-1M	QSFP+ to SFP+ 10-Gigabit Ethernet Direct Attach Copper (twinax copper cable) 1 m
QFX-QSFP-DACBO-3M	QSFP+ to SFP+ 10-Gigabit Ethernet Direct Attach Copper (twinax copper cable) 3 m
QFX-QSFP-DAC-1M	QSFP+ to QSFP+ Ethernet Direct Attach Copper (twinax copper cable) 1 m passive
QFX-QSFP-DAC-3M	QSFP+ to QSFP+ Ethernet Direct Attach Copper (twinax copper cable) 3 m passive

Model Number	Description
Software Licenses	
QFX3008-JSL-DRCTR-FAB	QFX3000-G base QFabric software
QFX3000M-JSL-DRCTR-FAB	QFX3000-M base QFabric software
QFX3000-JSL-EDGE-FAB	QFX3000 QFabric Switch QFabric Node feature license
QFX-JSL-DRCTR-ADV1	QFabric architecture advanced feature license for IS-IS, BGP, and IPv6 Routing
QFX-JSL-DRCTR-FC	QFabric architecture feature license for FCoE-to-FC gateway
QFX-JSL-DRCTR-FC-C16	QFabric architecture feature license for FCoE-to-FC gateway (capacity 16)
QFX-NODE-KIT	Conversion kit for QFX3500 Switch (standalone) to QFabric Node

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 2015 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.