

EX4600 Ethernet Switch



Product Overview

The EX4600 line of Ethernet switches delivers data center class scale, high availability, and high performance to campus distribution deployments. The EX4600 offers fixed 10 Gbps and modular 10 Gbps and 40 Gbps ports, giving it the flexibility to also be implemented in data center top-of-rack and service provider aggregation deployments.

Product Description

Featuring up to 72 wire-speed 1GbE or 10GbE small form-factor pluggable and pluggable plus transceiver (SFP/SFP+) ports, and up to 12 wire-speed 40GbE quad SFP+ transceiver (QSFP+) ports in a compact one rack unit (1 U) platform, the Juniper Networks® EX4600 Ethernet Switch delivers 1.44 Tbps of Layer 2 and Layer 3 connectivity to networked devices such as secure routers, servers, and other switches. The EX4600 base switch provides 24 fixed 10GbE SFP/SFP+ ports and 4 fixed 40GbE QSFP+ ports. A total of four models are available: two featuring AC power supplies and front-to-back or back-to-front airflow; and two featuring DC power supplies and front-to-back or back-to-front airflow.

All versions feature two expansion slots that can accommodate optional expansion modules, providing tremendous configuration and deployment flexibility for enterprise distribution networks. Two expansion modules are available:

- 8x10GBASE SFP/SFP+ fiber expansion module
- 4x40GbE QSFP+ expansion module¹

Architecture and Key Components

Virtual Chassis Technology

The EX4600 supports Juniper Networks' unique Virtual Chassis technology, which enables up to 10 interconnected switches to operate as a single, logical device with a single IP address. Virtual Chassis technology enables enterprises to separate physical topology from logical groupings of endpoints, ensuring efficient resource utilization. The EX4600 can participate in the same Virtual Chassis configuration with the Juniper Networks EX4300 Ethernet Switch, delivering highly flexible and scalable configuration options for enterprise distribution deployments.

EX4600 switches in a Virtual Chassis configuration can be connected using up to 16 of the 10GbE or 40GbE ports in link aggregation groups (LAGs) across 10GbE/40GbE ports, providing an aggregate backplane capacity of up to 1.04 Tbps.

In the data center, EX4600 Virtual Chassis deployments can extend across multiple top-of-rack or end-of-row switches, providing tremendous configuration flexibility for 10GbE server connectivity by only requiring redundant links between Virtual Chassis groups rather than each physical switch to ensure high availability. In addition, mixed Virtual Chassis configurations featuring EX4300 and EX4600 switches provide an ideal solution for data centers with a mix of 1GbE and 10GbE servers, or for environments transitioning from 1GbE to 10GbE server connectivity.

Enterprise Advantages

The EX4600 delivers a highly available, simple, and scalable 10GbE solution in a compact and power-efficient platform, making it a perfect fit for enterprise networks.

Simple

Virtual Chassis technology on the EX4600 dramatically simplifies the enterprise distribution layer, eliminating the need for Spanning Tree Protocol (STP), Virtual Router Redundancy Protocol (VRRP), complex routing, or VLAN configurations. In addition, when deployed in a Virtual Chassis configuration, the EX4600 reduces the number of devices to manage and can even eliminate network switching layers. A single Virtual Chassis configuration can span multiple nodes, reducing the need for larger, more expensive nodes on superior tiers. Juniper Networks Junos® operating system, which is common across the entire EX Series Ethernet Switches product family, accelerates the learning curve by enabling the entire architecture to be managed with a common set of instructions from a single pane of glass.

Scalable

A single EX4600 can support up to 72 10GbE ports (using 10GbE breakout cables on 40GbE fixed ports) at line rate, providing a highly scalable solution for the most demanding environments. In addition, Virtual Chassis technology makes it easy to scale the network while reducing management complexity. By adding switches to a Virtual Chassis configuration, it is possible to grow the number of switch ports without increasing the number of devices to manage. As switches are added to a Virtual Chassis

configuration, backplane bandwidth demands can be scaled to maintain adequate oversubscription ratios. Using just two 40GbE interfaces, Virtual Chassis bandwidth can start at 160 Gbps and grow to 960 Gbps, forming a LAG on all 12 40GbE interfaces.

Highly Available

When EX4600 switches are deployed in a Virtual Chassis configuration, Junos OS initiates a process to assign a master (active) and backup (hot standby) switch to manage the configuration. If the master switch fails, an integrated L2 and L3 graceful Routing Engine switchover (GRES) feature automatically switches to the backup to maintain continuous and uninterrupted system operations. The EX4600 switch also includes a number of other high availability features, including redundant power and cooling, graceful protocol restart, equal-cost multipath (ECMP), LAG links distributed across multiple slots, and quality of service (QoS) to prioritize time-sensitive data, nonstop bridging (NSB), nonstop active routing (NSR), and unified in-service software upgrades (unified ISSU). Unified ISSU is currently supported only in standalone configurations; support in Virtual Chassis configurations will be enabled in a future release.

Adaptable and Modular

As business grows, the network needs to evolve. With its compact 1 U form factor and complete switch features, the EX4600 can be easily redeployed in different environments and physical layouts. EX4600 switch capacity can be expanded without increasing the number of managed devices by simply adding more switches to a Virtual Chassis configuration.

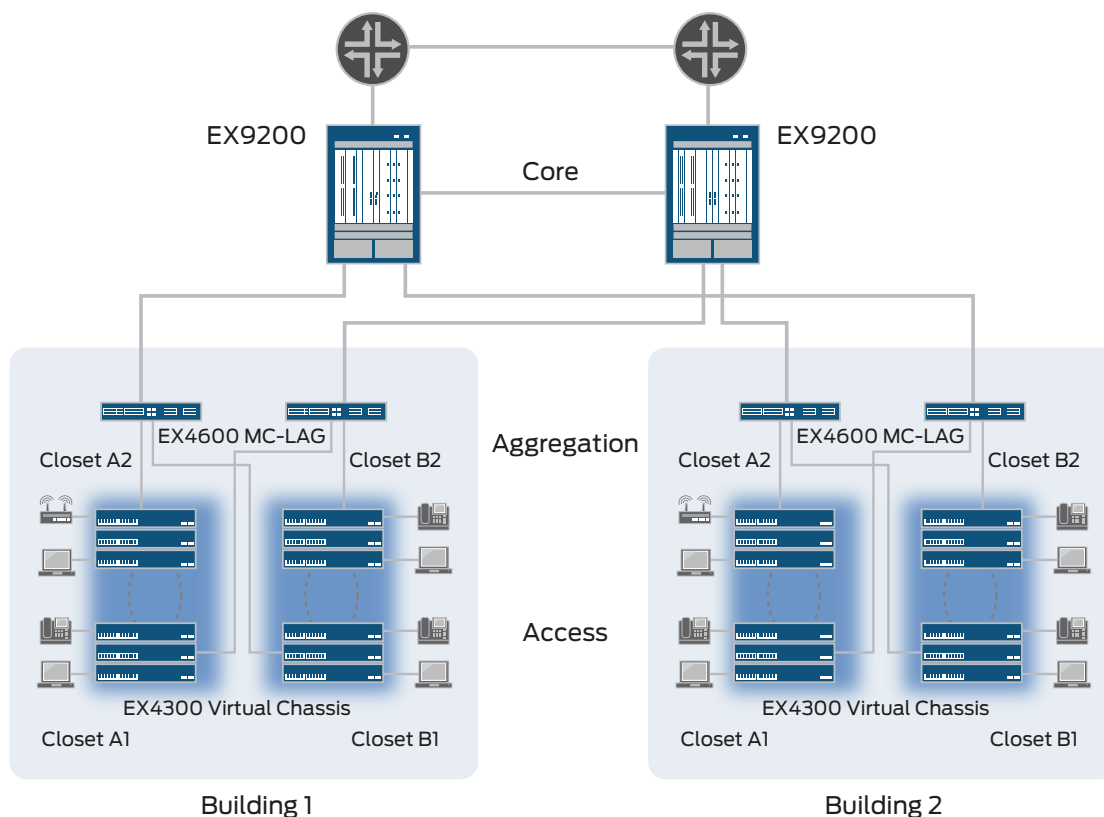


Figure 1: EX4600 as an enterprise distribution switch with MC-LAG

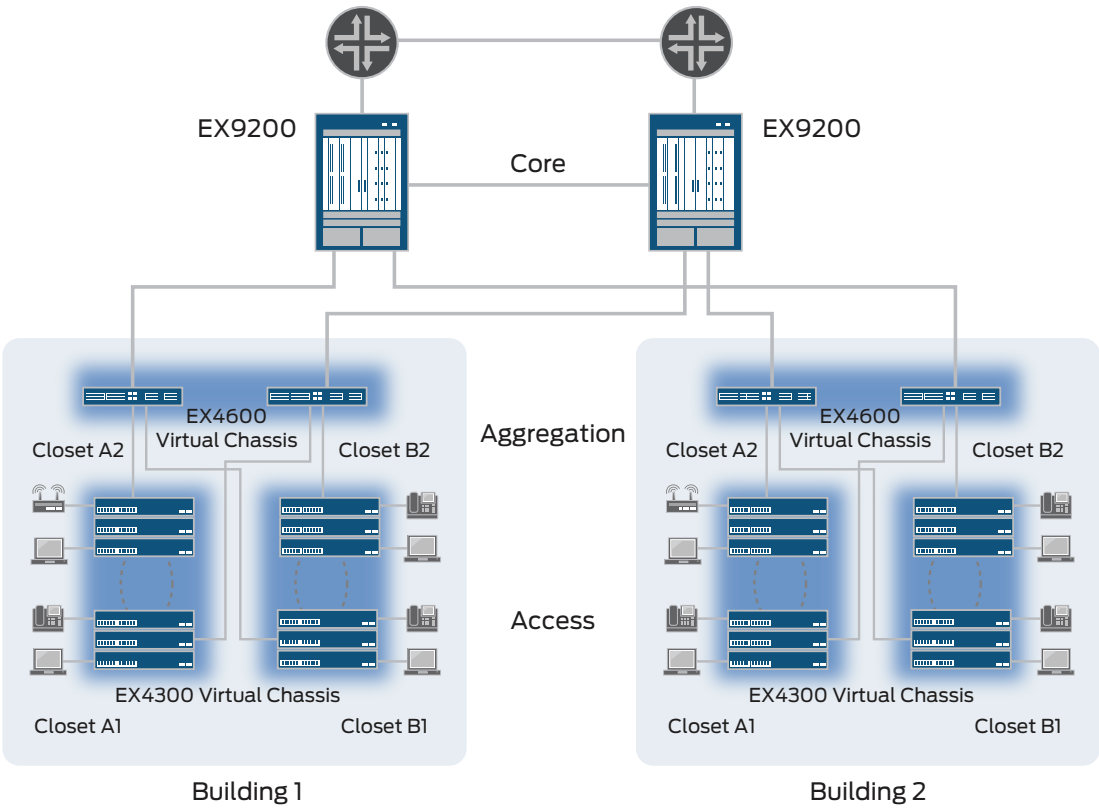


Figure 2: EX4600 as an enterprise distribution switch in a Virtual Chassis configuration

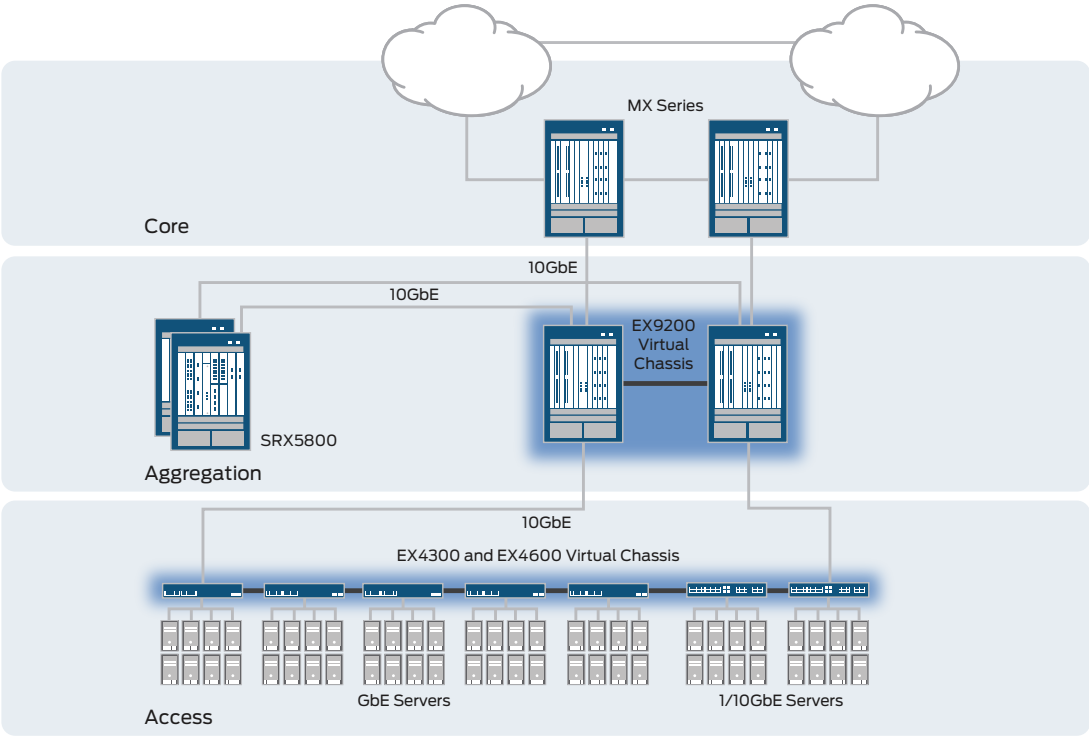


Figure 3: EX4600 provides 10GbE server access in the data center.

Enterprise Deployments

The EX4600 offers an economical, power-efficient, and compact solution for aggregating 10GbE expansions from access devices in building and enterprise deployments. The switch's dual-speed interfaces also support environments transitioning from 1GbE to 10GbE. The EX4600 can be deployed in the distribution layer with multichassis link aggregation (MC-LAG) (see Figure 1) to deliver higher resiliency with a distributed control plane, NSB, NSR, and unified ISSU. Multichassis LAG enables two EX4600 switches to act as separate devices with their own control planes, while eliminating STP by allowing link aggregation on the connected devices. In addition, unified ISSU allows each of the EX4600 switches to be upgraded individually without service interruption.

If manageability is a concern, the EX4600 can be deployed in a Virtual Chassis configuration to create a single management and control plane, NSB, NSR, and nonstop software upgrade (NSSU) (see Figure 2).

The EX4600 easily meets enterprise core switch requirements, delivering wire-speed performance on every port, full device redundancy, support for L3 dynamic routing protocols such as RIP and OSPF, L2 and L3 MPLS VPNs, and a comprehensive security and QoS feature set.

Data Center Deployments

The EX4600 Ethernet Switch is designed for low-density data center applications where high performance, high availability, and energy efficiency are key requirements (see Figure 3). Operating at wire speed, EX4600 switches deliver up to 1.07 Bpps throughput and a data rate of 1.44 Tbps (full duplex) for both L2

and L3 protocols. EX4600 switches can be interconnected in a Virtual Chassis configuration that also includes EX4300 switches, creating a single logical device that offers a variety of port and density options for mixed server environments.

Flexible deployment options enable the EX4600 to support back-to-front and front-to-back cooling, which ensures consistency with server designs for hot and cold aisle deployments. Front and rear facing configuration options ensure closer proximity to server ports, optimizing performance and keeping cable lengths short and manageable.

For small data centers, the EX4600, when deployed in a Virtual Chassis configuration, is ideally suited as an aggregation/core switch, aggregating 10GbE expansions from EX4300 Virtual Chassis configurations in the access layer (see Figure 4).

Customers introducing 10GbE into their server racks can use the EX4600 to add 10GbE-attached servers, iSCSI, and network-attached storage (NAS) with minimal impact to their current switching infrastructure.

10GbE Data Center Bridging and I/O Convergence Deployments

The EX4600 is a full IEEE DCB-based and T11 FC-BB-5-based Fibre Channel over Ethernet (FCoE) transit switch, delivering a high-performance solution for converged server edge access environments. As an FCoE transit switch, the EX4600 provides a pure IEEE DCB converged access layer between FCoE-enabled servers and an FCoE-enabled Fibre Channel (FC) storage area network (SAN) (see Figure 5).

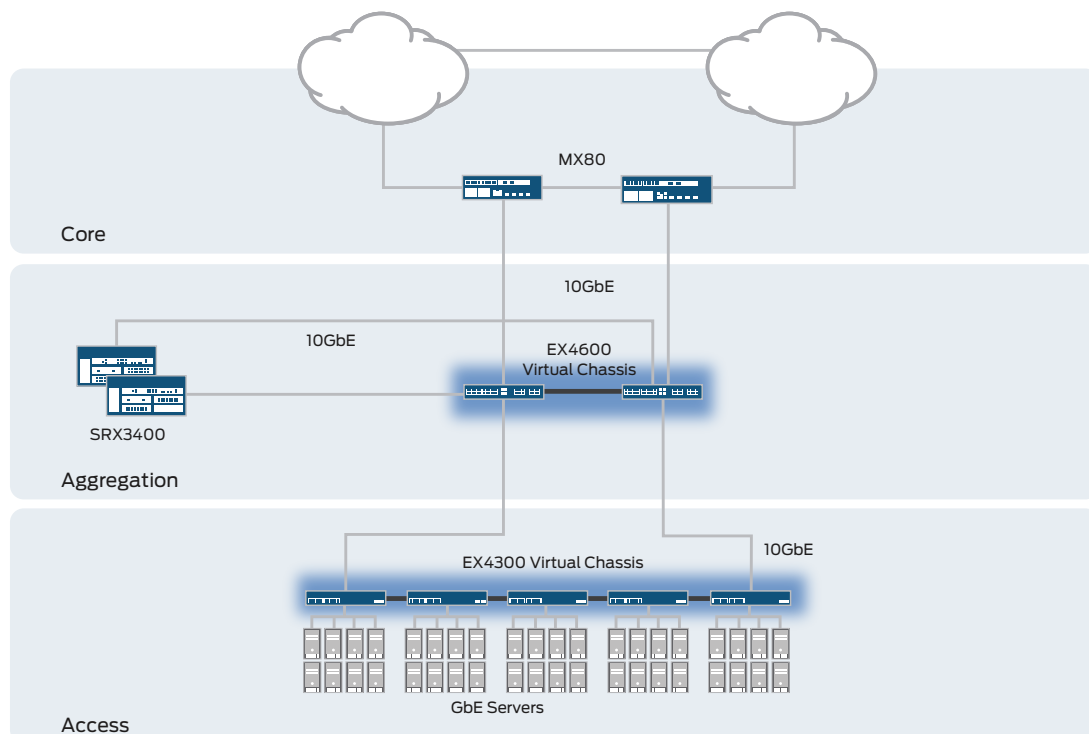


Figure 4: The EX4600 is ideal for small data center core deployments.

The EX4600 also supports FCoE Initialization Protocol (FIP) snooping, which provides perimeter protection to ensure that the presence of an Ethernet layer does not impact existing SAN security policies. The FCoE transit switch functionality, along with priority-based flow control (PFC) and Data Center Bridging Capability Exchange protocol (DCBX), are included as part of the default software; no additional licenses are required.

Junos Operating System

The EX4600 runs the same Juniper Networks Junos operating system that is used by other EX Series Ethernet Switches, as well as all 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 all products. To maintain that consistency, Junos OS adheres to a highly disciplined development process that uses a single source code, follows a single quarterly release train, 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 all Junos OS-powered products to be updated simultaneously with the same software release. All features are fully regression tested, making each new release a true superset of the previous version. Customers can deploy the software with complete confidence that all existing capabilities will be maintained and operate in the same way.

Management Options

Four different system management options are available for the EX4600 Ethernet Switch:

- The standard Junos OS CLI offers the same granular management capabilities and scripting parameters found in any router powered by the Junos operating system.
- The EX4600 also includes the integrated Juniper Networks Junos Web Software, an embedded web-based device manager that allows users to configure, monitor, troubleshoot, and perform system maintenance on individual switches via a browser-based graphical interface.
- EX4600 performance, configuration, and fault 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 EX4600 is also supported by Juniper Networks Junos Space Network Director, a smart, comprehensive, and automated network management tool that enables network administrators to visualize, analyze, and control their entire enterprise network—data center and campus, physical and virtual, wired and wireless—through a single pane of glass.

Junos Space is 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. 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.

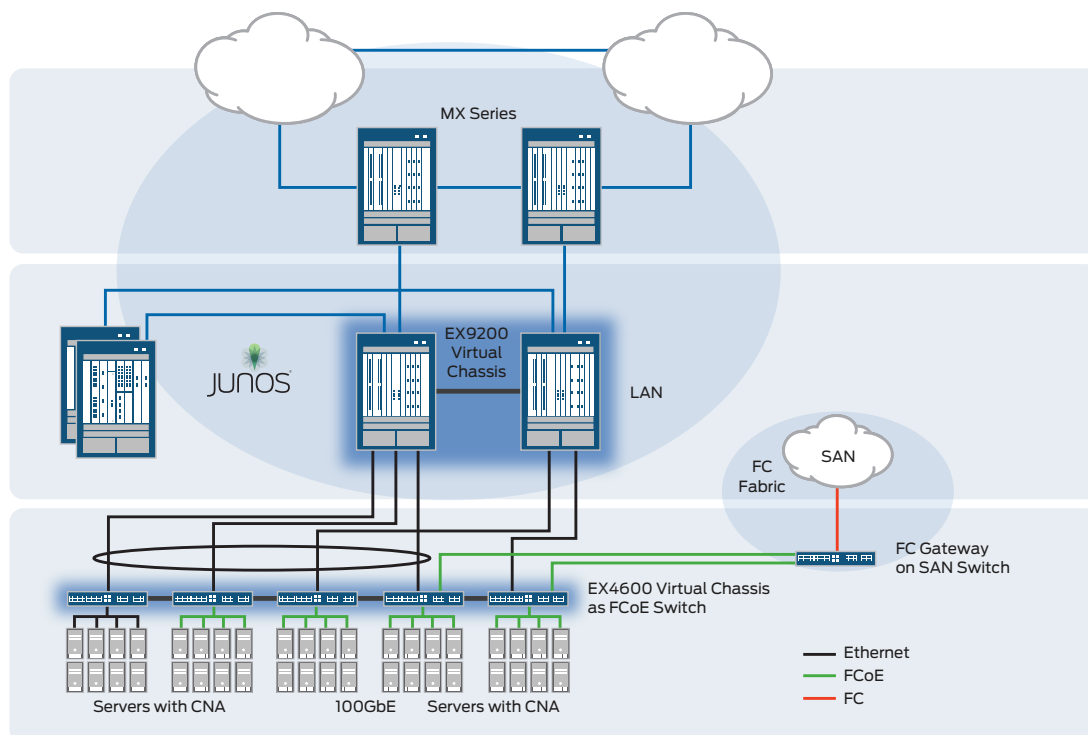


Figure 5: The EX4600 can be deployed as a data center FCoE transit switch in a top-of-rack Virtual Chassis configuration.

Features and Benefits


EX4600 Ethernet switches include the following key features and benefits:

- **High performance:** Each EX4600 supports up to 1.44 Tbps of bandwidth or 1.07 Bpps at the minimum Ethernet frame size, with 24 line-rate 10GbE fixed ports, 4 40GbE fixed ports, and 2 expansion slots that can each support either a 4x40GbE module or an 8x10GbE module.
- **Unified in-service software upgrade (unified ISSU):** With its Intel core processor, the EX4600 switch allows Junos OS to run within a virtual machine (VM) on Linux. Junos OS runs in two separate VMs in active and standby pairs; during software upgrade cycles, the switches seamlessly move to the newer software version while maintaining intact data plane traffic. This true topology-independent ISSU (TISSU), an industry-first software upgrade feature for a fixed-configuration top-of-rack switch, is supported across all L2 and L3 protocols and doesn't need the support of any other switches to perform an image upgrade.
- **High availability:** EX4600 switches offer dual internal load-sharing AC power supplies and redundant variable-speed fans as standard features, protecting the switch from a single power supply or fan failure. DC power options are also available.
- **Automation:** The EX4600 switches support a number of features for network automation and plug-and-play operations. Features include zero touch provisioning (ZTP), operations and event scripts, automatic rollback, and Python scripting. The switch also offers support for integration with VMware NSX Layer 2 Gateway Services, Puppet, and OpenStack.
- **Energy efficient:** Consuming less than five watts per 10GbE interface, the EX4600 offers a low power solution for 10GbE top-of-rack, end-of-row, and distribution deployments. The EX4600 switches also improve cooling efficiency with redundant variable-speed fans that automatically adjust their speed based on existing conditions to reduce power consumption.
- **Small footprint:** The EX4600 supports up to 72 wire-speed 10GbE ports in a single 1 U platform.
- **Flexible forwarding table:** The EX4600's flexible forwarding table allows the hardware table to be carved into configurable partitions of L2 media access control (MAC), L3 host, and longest prefix match (LPM) tables. In a pure Layer 2 environment, the EX4600 supports up to 288,000 MAC addresses. In Layer 3 mode, the table can support up to 128,000 host entries; in LPM mode, it can support up to 128,000 prefixes. Junos OS provides configurable options through a CLI so that each EX4600 can be optimized for different deployment scenarios.
- **Intelligent buffer management:** EX4600 switches have a total of 12 MB shared buffers. While 25 percent of the total buffer space is dedicated, the rest is shared among all ports and is user configurable. The intelligent buffer mechanism in the EX4600 effectively absorbs traffic bursts while providing deterministic performance, significantly increasing performance over static allocation.
- **Insight technology for analytics:** The EX4600 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), comma-separated values (CSV), and tab-separated values (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:** EX4600 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 EX4600 can be deployed as a low-latency MPLS label-switching router (LSR) or MPLS provider edge (PE) router in smaller scale environments. The EX4600 is the industry's only compact, low-latency, high-density, low-power switch to offer an MPLS feature set.
- **MACsec:** The EX4600 is capable of MACsec features on all 10GbE ports to support 400 Gbps of near line-rate hardware-based traffic encryption on all fiber ports, including the base unit and optional 10GbE expansion modules. MACsec in software will be enabled in a future release.

Additional Features

- System status LEDs
- Versatile two- and four-post rack mounting options
- Front-to-back and back-to-front airflow options
- AC and DC power supply options
- Support for jumbo frames (9,000)
- Quality of service (IEEE 802.1p marking)
- Multicast (Internet Group Management Protocol v1/v2/v3 snooping)
- Layer 2 features including support for 4,096 VLAN IDs, Spanning Tree (802.1s and 802.1w), bridge protocol data unit (BPDU) guard, 802.3as Link Distribution
- Management features including Telnet and SSH v1/v2, SNMP v1-v3, RADIUS, TACACS+, and RMON

SKU	Configuration	Power Consumption
EX4600-40F-AFO	<ul style="list-style-type: none"> Fixed 10GbE ports with 10G-USR optics, all ports forwarding (line rate), <10 m 4 fixed 40GbE ports with 40G-SR4 optics 1 4x40GbE QIC card with 4 40G-SR4 optics 1 8x10GbE QIC card with 8 10G-USR optics, all ports forwarding (line rate), <10 m Dual power supplies 	• 279 W
EX4600-40F-AFI		• 285 W
EX4600-40F-DC-AFO		• 392 W
EX4600-40F-DC-AFI		• 320 W



EX4600 Switch Specifications

Hardware

- Switching capacity: 1.44 Tbps/1.07 Bpps
- Weight: 21.4 lb (9.9 kg)
- Dimensions (HxWxD): 1.72 x 17.36 x 20.48 in (4.37 x 44.09 x 52.02 cm)
- Switching mode: Cut-through and store-and-forward
- Front-to-back or back-to-front airflow (for hot aisle/cold aisle deployment)
- Management and rear console port connections
- Predicted mean time between failures (MTBF): 150,000 hours
- Predicted FIT rate: 4,987

Interface Options

- 1GbE SFP: 24(40) (with 10GbE expansion modules)
- 10GbE SFP+: 24(40/72) (with 10GbE expansion modules/with fixed 40GbE ports using breakout cables)
- 40GbE QSFP+: 4(12) (with expansion modules)
 - Each fixed QSFP+ port can be configured as a 4x10GbE interface
 - Each QSFP+ port can be configured as a 40 Gbps port
 - USB port
 - Console port
 - 2 management ports: 1 RJ-45 and 1 SFP
 - Supported transceiver and direct attach cable
 - SFP+ 10GbE optical modules
 - SFP+ DAC cables: 1/3/5 m direct-attached copper and 1/3/5/7/10 m active direct-attached copper

- SFP GbE optical and copper module
- QSFP+ to SFP+ 10GbE direct attach break-out copper (1/3 m direct-attached copper cable)

Rack Installation Kit

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

Airflow

- Front-to-back and back-to-front cooling
- Redundant variable-speed fans to reduce power draw
- Power Supply and Fan Modules
- Dual redundant (1+1) and hot-pluggable power supplies
- 110-240 V single phase AC power
- 36 to -72 V DC power
- Redundant (N+1) and hot-pluggable fan modules for front-to-back and back-to-front airflow
- Performance Scale (Unidimensional)
- MAC addresses per system: 288,000*
- VLAN IDs: 4,091
- Number of LAGs: 128
- Number of ports per LAG: 32
- FCoE scale:
 - Number of FCoE VLANs/FC virtual fabrics: 4,095
- Firewall filters: 4,000
- IPv4 unicast routes: 128,000 prefixes; 208,000 host routes; 64** ECMP paths
- IPv4 multicast routes: 104,000
- IPv6 multicast routes: 52,000
- IPv6 unicast routes: 64,000 prefixes
- Address Resolution Protocol (ARP) entries: 48,000
- Jumbo frame: 9,216 bytes

Spanning Tree Protocol (STP)

- Multiple Spanning Tree Protocol (MSTP) instances: 64
- VLAN Spanning Tree Protocol (VSTP) instances: 253

Traffic Mirroring

- Mirroring destination ports per switch: 4
- Maximum number of mirroring sessions: 4
- Mirroring destination VLANs per switch: 4

Layer 2 Features

- STP—IEEE 802.1D (802.1D-2004)
- Rapid Spanning Tree Protocol (RSTP) (IEEE 802.1w); MSTP (IEEE 802.1s)
- Bridge protocol data unit (BPDU) protect
- Loop protect
- Root protect
- RSTP and VSTP running concurrently
- VLAN—IEEE 802.1Q VLAN trunking
- Routed VLAN interface (RVI)
- Port-based VLAN
- MAC address filtering
- Layer 2 protocol tunneling (L2PT)
- QinQ**
- VLAN translation

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

** Roadmap

- Static MAC address assignment for interface
- Per VLAN MAC learning (limit)
- MAC learning disable
- Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)

Link Aggregation

- Multichassis link aggregation (MC-LAG)
- Redundant trunk group (RTG)
- LAG load sharing algorithm—bridged or routed (unicast or multicast) traffic:
 - IP: SIP, Dynamic Internet Protocol (DIP), TCP/UDP source port, TCP/UDP destination port
 - Layer 2 and non-IP: MAC SA, MAC DA, Ethertype, VLAN ID, source port
 - FCoE packet: Source ID (SID), destination ID (DID), originator exchange ID (OXID), source port

Layer 3 Features (IPv4)

- Static routing
- Routing policy
- Routing protocols (RIP, OSPF, IS-IS, BGP)
- Virtual Router Redundancy Protocol (VRRP)
- Bidirectional Forwarding Detection (BFD) protocol
- Virtual router
- Dynamic Host Configuration Protocol (DHCP) relay
- Proxy Address Resolution Protocol (ARP)
- Multicast Features
- Internet Group Management Protocol (IGMP): v1, v2, v3
- IGMP snooping: v1, v2, v3
- IGMP filter
- PIM-SM
- Multicast Source Discovery Protocol (MSDP)
- Security and Filters
- Secure interface login and password
- RADIUS
- TACACS+
- Ingress and egress filters: Allow and deny, port filters, VLAN filters, and routed filters, including management port filters
- Filter actions: Logging, system logging, reject, mirror to an interface, counters, assign forwarding class, permit, drop, police, mark
- SSH v1, v2
- Static ARP support
- Storm control, port error disable, and autorecovery
- Control plane denial-of-service (DoS) protection
- IP source guard
- Dynamic ARP inspection (DAI)
- Sticky MAC address
- DHCP snooping

Quality of Service (QoS)

- L2 and L3 QoS: Classification, rewrite, queuing
- Rate limiting:
 - Ingress policing: 1 rate 2 color, 2 rate 3 color
 - Egress policing: Policer, policer mark down action
 - Egress shaping: Per queue, per port
- 12 hardware queues per port (8 unicast and 4 multicast)
- Strict priority queuing (LLQ), smoothed deficit weighted round-robin (SDWRR), weighted random early detection (WRED), weighted tail drop
- 802.1p remarking
- L2 classification criteria: Interface, MAC address, Ethertype, 802.1p, VLAN
- Congestion avoidance capabilities: WRED
- Trust IEEE 802.1p (ingress)
- Remarking of bridged packets
- Priority-based flow control (PFC)—IEEE 802.1Qbb
- Enhanced Transmission Selection (ETS)—IEEE 802.1Qaz
- Data Center Bridging Exchange Protocol (DCBX), DCBx FCoE, and iSCSI type, length, and value (TLVs)
- Fibre Channel over Ethernet (FCoE)
- FCoE transit switch (FIP snooping ACL installation)
- Virtual fiber channel gateway
- FCoE session path learning
- FCoE session health monitoring
- Graceful restart for FIP snooping
- FC-BB-6 VN2VN snooping

Virtual Chassis

- 40GbE and 10GbE as Virtual Chassis port
- Virtual Chassis Routing Engine (RE) election
- Virtual Chassis pre-provisioning (plug and play)
- Auto-LAG formation of Virtual Chassis ports
- Mixed Virtual Chassis support between EX4300-EX4600 (in data center only)
- FCoE transit across Virtual Chassis members
- QoS on Virtual Chassis ports
- Local designated forwarding
- Graceful RE switchover (GRES)
- Nonstop routing (NSR)
- Nonstop bridging (NSB)
- Monitor distributed aggregate interface
- Control plane protection for virtual RE

High Availability

- Unified ISSU
- Bidirectional Forwarding Detection (BFD)
- Uplink failure detection (UFD)

MPLS

- Static label-switched paths (LSPs)
- RSVP-based signaling of LSPs
- LDP-based signaling of LSPs
- LDP tunneling (LDP over RSVP)
- MPLS class of service (CoS)
- MPLS access control list (ACL)/policers
- MPLS LSR support
- IPv6 tunneling (6PE) (via IPv4 MPLS backbone)
- MPLS Operation, Administration, and Maintenance (OAM)—
- LSP ping
- IPv4 L3 VPN (RFC 2547, 4364)

Management and Operations

- Junos Space Network Director**
- Role-based CLI management and access
- CLI via console, telnet, or SSH
- Extended ping and traceroute
- Junos OS configuration rescue and rollback
- Image rollback
- SNMP v1/v2/v3
- Junos XML management protocol
- sFlow v5
- High frequency statistics collection
- Beacon LED for port and system
- Automation and orchestration
- Zero touch provisioning (ZTP)
- OpenStack Neutron Plug-in
- Puppet
- Chef
- Python
- Junos OS event, commit, and operations scripts

Traffic Mirroring

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

Standards Compliance

IEEE Standards

- IEEE 802.1D
- IEEE 802.1w
- IEEE 802.1
- IEEE 802.1Q

- IEEE 802.1p
- IEEE 802.1ad
- IEEE 802.3ad
- IEEE 802.1AB
- IEEE 802.3x
- IEEE 802.1Qbb
- IEEE 802.1Qaz
- IEEE 802.1Qau**
- IEEE 802.1Qbg**

Supported RFCs

- RFC 768 UDP
- RFC 783 Trivial File Transfer Protocol (TFTP)
- RFC 791 IP
- RFC 792 ICMP
- RFC 793 TCP
- RFC 826 ARP
- RFC 854 Telnet client and server
- RFC 894 IP over Ethernet
- RFC 903 RARP
- RFC 906 TFTP Bootstrap
- RFC 951 1542 BootP
- RFC 1058 Routing Information Protocol
- RFC 1112 IGMP v1
- RFC 1122 Host requirements
- RFC 1142 OSI IS-IS Intra-domain Routing Protocol
- RFC 1256 IPv4 ICMP Router Discovery (IRDP)
- RFC 1492 TACACS+
- RFC 1519 Classless Interdomain Routing (CIDR)
- RFC 1587 OSPF not-so-stubby area (NSSA) Option
- RFC 1591 Domain Name System (DNS)
- RFC 1745 BGP4/IDRP for IP—OSPF Interaction
- RFC 1772 Application of the Border Gateway Protocol in the Internet
- RFC 1812 Requirements for IP Version 4 routers
- RFC 1997 BGP Communities Attribute
- RFC 2030 SNTP, Simple Network Time Protocol
- RFC 2068 HTTP server
- RFC 2131 BOOTP/DHCP relay agent and Dynamic Host
- RFC 2138 RADIUS Authentication
- RFC 2139 RADIUS Accounting
- RFC 2154 OSPF w/Digital Signatures (Password, MD-5)
- RFC 2236 IGMP v2
- RFC 2267 Network ingress filtering
- RFC 2328 OSPF v2 (edge mode)
- RFC 2338 VRRP
- RFC 2362 PIM-SM (edge mode)
- RFC 2370 OSPF Opaque link-state advertisement (LSA) Option

** Roadmap

- RFC 2385 Protection of BGP Sessions via the TCP Message Digest 5 (MD5) Signature Option
- RFC 2439 BGP Route Flap Damping
- RFC 2453 RIP v2
- RFC 2474 Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers
- RFC 2597 Assured Forwarding PHB (per-hop behavior) Group
- RFC 2598 An Expedited Forwarding PHB
- RFC 2697 A Single Rate Three Color Marker
- RFC 2698 A Two Rate Three Color Marker
- RFC 2796 BGP Route Reflection—An Alternative to Full Mesh IBGP
- RFC 2918 Route Refresh Capability for BGP-4
- RFC 3065 Autonomous System Confederations for BGP
- RFC 3376 IGMP v3 (source-specific multicast include mode only)
- RFC 3392 Capabilities Advertisement with BGP-4
- RFC 3446 Anycast RP
- RFC 3569 SSM
- RFC 3618 MSDP
- RFC 3623 Graceful OSPF Restart
- RFC 4271 Border Gateway Protocol 4 (BGP-4)
- RFC 4360 BGP Extended Communities Attribute
- RFC 4456 BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)
- RFC 4486 Subcodes for BGP Cease Notification Message
- RFC 4724 Graceful Restart Mechanism for BGP
- RFC 4812 OSPF Restart Signaling
- RFC 4893 BGP Support for Four-octet AS Number Space
- RFC 5176 Dynamic Authorization Extensions to RADIUS
- RFC 5396 Textual Representation of Autonomous System (AS) Numbers
- RFC 5668 4-Octet AS Specific BGP Extended Community
- RFC 5880 Bidirectional Forwarding Detection (BFD)
- Dynamic Host Configuration Protocol (DHCP) server

Supported MIBs

- RFC 1155 SMI
- RFC 1157 SNMPv1
- RFC 1212, RFC 1213, RFC 1215 MIB-II, Ethernet-Like MIB, and TRAPs
- RFC 1850 OSPFv2 MIB
- RFC 1901 Introduction to Community-based SNMPv2
- RFC 2011 SNMPv2 for Internet protocol using SMIv2
- RFC 2012 SNMPv2 for transmission control protocol using SMIv2
- RFC 2013 SNMPv2 for user datagram protocol using SMIv2
- RFC 2233, The Interfaces Group MIB using SMIv2
- RFC 2287 System Application Packages MIB
- RFC 2570 Introduction to Version 3 of the Internet-standard

Network Management Framework

- RFC 2571 An Architecture for describing SNMP Management Frameworks (read-only access)
- RFC 2572 Message Processing and Dispatching for the SNMP (read-only access)
- RFC 2576 Coexistence between SNMP Version 1, Version 2, and Version 3
- RFC 2578 SNMP Structure of Management Information MIB
- RFC 2579 SNMP Textual Conventions for SMIv2
- RFC 2580 Conformance Statements for SMIv2
- RFC 2665 Ethernet-like interface MIB
- RFC 2787 VRRP MIB
- RFC 2790 Host Resources MIB
- RFC 2819 RMON MIB
- RFC 2863 Interface Group MIB
- RFC 2932 IPv4 Multicast MIB
- RFC 3410 Introduction and Applicability Statements for Internet Standard Management Framework
- RFC 3411 An architecture for describing SNMP Management Frameworks
- RFC 3412 Message Processing and Dispatching for the SNMP
- RFC 3413 Simple Network Management Protocol (SNMP) (all MIBs are supported except the Proxy MIB)
- RFC 3414 User-based Security Model (USM) for version 3 of SNMPv3
- RFC 3415 View-based Access Control Model (VACM) for the SNMP
- RFC 3416 Version 2 of the Protocol Operations for the SNMP
- RFC 3417 Transport Mappings for the SNMP
- RFC 3418 Management Information Base (MIB) for the SNMP
- RFC 3584 Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework
- RFC 3826 The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model
- RFC 4188 Definitions of Managed Objects for Bridges
- RFC 4318 Definitions of Managed Objects for Bridges with Rapid Spanning Tree Protocol
- RFC 4363b Q-Bridge VLAN MIB

Approvals

Safety

- CAN/CSA-C22.2 No. 60950-1 (2007) Information Technology Equipment—Safety
- UL 60950-1 (2nd Ed.) Information Technology Equipment—Safety
- EN 60950-1 (2005) Information Technology Equipment—Safety
- IEC 60950-1 (2005) Information Technology Equipment—Safety (All country deviations): CB Scheme report.
- EN 60825-1 +A1+A2 (1994) Safety of Laser Products—Part 1: Equipment Classification

- GR-63-Core (2006) Network Equipment, Building Systems (NEBS) Physical Protection
- GR-1089-Core (2006) EMC and Electrical Safety for Network Telecommunications Equipment
- SR-3580 (1995) NEBS Criteria Levels (Level 3)

EMC

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

Environmental Compliance



Restriction of Hazardous Substances (ROHS) 6/6



China Restriction of Hazardous Substances (ROHS)



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



Waste Electronics and Electrical Equipment (WEEE)



Recycled material



80 Plus Silver PSU Efficiency

Telco

- Common Language Equipment Identifier (CLEI) code

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 nonoperating: 0% to 95% (noncondensing)

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
Switch Hardware	
EX4600-40F-AFO	24 SFP+/SFP ports, 4 QSFP+ ports, 2 expansion slots, redundant fans, 2 AC power supplies, front to back airflow
EX4600-40F-AFI	24 SFP+/SFP ports, 4 QSFP+ ports, 2 expansion slots, redundant fans, 2 AC power supplies, back to front airflow
EX4600-40F-DC-AFO	24 SFP+/SFP ports, 4 QSFP+ ports, 2 expansion slots, redundant fans, 2 DC power supplies, front to back airflow
EX4600-40F-DC-AFI	24 SFP+/SFP ports, 4 QSFP+ ports, 2 expansion slots, redundant fans, 2 DC power supplies, back to front airflow
Modules and Spares	
QFX-EM-4Q	4-port QSFP+ expansion module for EX4600-24Q
EX4600-EM-8F	8-port SPF+/SFP expansion module for EX4600 and QFX5100-24Q
JPSU-650W-AC-AFO	AC 650 W PSU, front-to-back airflow for EX4600-48S, EX4600-24Q
JPSU-650W-AC-AFI	AC 650 W PSU, back-to-front airflow for EX4600-48S, EX4600-24Q
JPSU-650W-DC-AFO	DC 650 W PSU, front-to-back airflow for EX4600-48S, EX4600-24Q
JPSU-650W-DC-AFI	DC 650 W PSU, back-to-front airflow for EX4600-48S, EX4600-24Q
QFX5100-FAN-AFI	Back-to-front airflow fan module for EX4600
QFX5100-FAN-AFO	Front-to-back airflow fan module for EX4600
QFX5100-EM-BLNK	Blank module for EX4600
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, 1,310 nm for 10 km transmission on single mode fiber-optic (SMF)
QFX-SFP-10GE-ER	SFP+ 10GBASE-ER 10 Gigabit Ethernet Optics, 1,550 nm for 40 km transmission on single-mode fiber (SMF)
EX-SFP-10GE-ZR	SFP+ 10GBASE-ZR 10 Gigabit Ethernet Optics, 1,550 nm for 80 km transmission on single-mode fiber (SMF)
QFX-SFP-DAC-1M	SFP+ 10 Gigabit Ethernet Direct Attach Copper (direct-attached copper cable) 1 m
QFX-SFP-DAC-3M	SFP+ 10 Gigabit Ethernet Direct Attach Copper (direct-attached copper cable) 3 m
QFX-SFP-DAC-5M	SFP+ 10 Gigabit Ethernet Direct Attach Copper (direct-attached copper cable) 5 m
QFX-SFP-DAC-1MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active direct-attached copper cable) 1 m
QFX-SFP-DAC-3MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active direct-attached copper cable) 3 m

Model Number	Description
QFX-SFP-DAC-5MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active direct-attached copper cable) 5 m
QFX-SFP-DAC-7MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active direct-attached copper cable) 7 m
QFX-SFP-DAC-10MA	SFP+ 10 Gigabit Ethernet Direct Attach Copper (active direct-attached copper cable) 10 m
QFX-QSFP-DAC-1M	QSFP+ to QSFP+ Ethernet Direct Attach Copper (direct-attached copper cable) 1 m passive
QFX-QSFP-DAC-3M	QSFP+ to QSFP+ Ethernet Direct Attach Copper (direct-attached copper cable) 3 m passive
JNP-QSFP-DAC-5M	QSFP+ to QSFP+ Ethernet Direct Attach Copper (direct-attached copper cable) 5 m passive
QFX-QSFP-DACBO-1M	QSFP+ to SFP+ 10 Gigabit Ethernet Direct Attach Breakout Copper (direct-attached copper cable) 1 m
QFX-QSFP-DACBO-3M	QSFP+ to SFP+ 10 Gigabit Ethernet Direct Attach Breakout Copper (direct-attached copper cable) 3 m
QFX-QSFP-40G-SR4	QSFP+ 40GBASE-SR4 40 Gigabit Optics, 850 nm for up to 150 m transmission on MMF
QFX-QSFP-40G-ESR4	QSFP+ 40GBASE-ESR4 40 Gigabit Optics, 300 m (400 m) with OM3 (OM4) MMF
JNP-QSFP-40G-LR4	QSFP+ 40GBASE-LR4 40 Gigabit Optics, 1,310 nm for up to 10 km transmission on SMF
QFX-SFP-1GE-T	SFP 1000BASE-T Copper Transceiver Module for up to 100 m transmission on Cat5
QFX-SFP-1GE-SX	SFP 1000BASE-SX Gigabit Ethernet Optics, 850 nm for up to 550 m transmission on MMF
QFX-SFP-1GE-LX	SFP 1000BASE-LX Gigabit Ethernet Optics, 1,310 nm for 10 km transmission on SMF

Advanced Feature Licenses

EX4600-AFL	EX4600 Advanced feature license for IS-IS, BGP, and MPLS
------------	--

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 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.