


QFX10002 Fixed Ethernet Switches



Product Overview

The QFX10000 line of high-performance fixed core and spine switches delivers industry-leading scalability, density, and flexibility, helping cloud and data center operators build automated data center networks that provide superior long-term investment protection. Designed for a diverse set of deployment options, the QFX10000 switches allow data center operators to build cloud networks that best suit their deployment needs and easily evolve as requirements change over time.

Product Description

Juniper Networks® QFX10002 line of fixed configuration switches offers 72-port and 36-port 40GbE options in a 2 U form factor, with support for 10GbE and 100GbE. Both switches support quad small form-factor pluggable plus transceiver (QSFP+) and QSFP28 ports for 40GbE and 100GbE speeds, respectively. With support for 10GbE, 40GbE, and 100GbE in a single platform, the QFX10002 switches provide the foundation for today's data center, delivering long-term investment protection for current and emerging requirements.

Built on custom, Juniper-designed Q5 ASICs, the high-performance, highly versatile QFX10002 switches scale from 2.88 Tbps to 5.76 Tbps of throughput. Additionally, by offering the industry's highest 100GbE port density—up to 24 in a 2 U fixed form factor—the QFX10002 enables a simple, consolidated network design and allows customers to seamlessly evolve data center networks from existing 10GbE and 40GbE architectures to 100GbE to keep pace with evolving needs.

The QFX10002 is a member of the QFX10000 line of Ethernet switches, which can be deployed in a number of different network designs and fabrics, including Layer 3 fabric and multichassis link aggregation (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 their pace of innovation.

QFX10000 switches are 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 switches support Data Center Interconnect (DCI), critical to workload mobility and application availability.

For networks evolving to SDN, the QFX10000 switches can integrate with OpenContrail and VMware NSX SDN controllers, giving users options for SDN systems. The QFX10000 performs both Virtual Extensible LAN (VXLAN) L2 and L3 gateway functionality; the open, standards-based platform also interoperates with Open vSwitch Database (OVSDB) to support automated management and control capabilities.

QFX10002 Highlights

- Up to 72 40GbE ports in a 2 U form factor
- Up to 5.76 Tbps Layer 2 and Layer 3 performance
- Unparalleled investment protection with support for 10GbE, 40GbE, and 100GbE in the same platform
- Highest logical L2/L3 scale—up to 512,000 media access control (MAC) addresses, 2 million host routes, and 2 million forwarding information base forwarding tables
- Deep buffers with up to 100 ms packet buffering per port
- No head-of-line blocking with Virtual Output Queue (VOQ)-based architecture



- Flexible network architectures including L3 fabric and Juniper's MC-LAG for L2 and L3 networks
- Juniper virtualized open network operating system framework for programmability through APIs
- Advanced Junos OS features such as BGP add-path, VXLAN routing, MPLS, and Fibre Channel over Ethernet (FCoE)
- Rich automation capabilities with Python, Chef Puppet, and zero touch provisioning (ZTP)

Architecture and Key Components

QFX10002 Switch Models

The QFX10002 line includes two fixed 2 U switch models, each providing wire-speed packet performance in 10GbE, 40GbE, or 100GbE speeds. Both switches provide a rich set of Junos OS features. In addition to a high throughput Q5 ASIC, the control plane performance on all QFX10002 switches is further enhanced with a powerful 2.5 GHz quad core Intel CPU and 16 GB of SDRAM.

- **QFX10002-36Q:** The QFX10002-36Q 40GbE switch features 36 QSFP+ ports or 12 QSFP28 100GbE ports with two redundant 1600 W AC/DC power supplies and three fan trays.
- **QFX10002-72Q:** The QFX10002-72Q 40GbE switch offers 72 QSFP+ ports or 24 QSFP28 100GbE ports with four redundant 1600 W AC/DC power supplies and three fan trays.

The Q5 ASIC



Figure 1: The custom-built Juniper Q5 ASIC

The QFX10002 switches are built with Juniper custom Q5 ASICs, which offer industry-leading performance and scale with 1 Tbps of switching throughput and support for network virtualization with VXLAN, Ethernet VPN (EVPN), and MPLS. The Q5 ASIC is 400GbE-ready and embedded with an on-chip analytics capability, along with a precision timing protocol and high-frequency monitoring.

Flexible Network Architectures

QFX10000 switches support a diverse set of deployment options, including L3 fabric and L2 and L3 networks with MC-LAG. Customers can choose the architecture that best suits their deployment needs and easily adapt and evolve as requirements change over time. The QFX10002 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 nonblocking and predictable performance and scale characteristics. For example, a two-tier fabric with QFX10000 switches as the spine and QFX5100 switches as leafs can scale to support more than 6,900 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 autonomous system (AS) numbers, routing policy, loopback address assignments, and many other implementation details. The second largest problem is automating the creation of an IP fabric at a large scale. Juniper has created the OpenClos project to provide free and open-source tools to automate the creation of IP fabrics in the data center. OpenClos is a set of Python scripts that has been developed as an open-source project and lives on GitHub. OpenClos takes a set of inputs that describe the shape and size of a data center and it outputs network switch configuration files and even a cabling plan.

- **MC-LAG:** QFX10000 switches eliminate Spanning Tree Protocol (STP) in traditional L2 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 deploy overlay networks to provide Layer 2 adjacencies for applications over Layer 3 fabrics. The overlay networks utilize VxLAN in the data plan and EVPN or OVSDB for programming the overlays. The overlays can operate without a controller or can be orchestrated with an SDN controller like OpenContrail or VMware NSX.

Carrier-Class Operating System

The QFX10000 switches run 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 switches 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, which ensures error-free network configurations
- A powerful set of scripts for on-box problem detection, reporting, and resolution




Figure 2: QFX10000 switches can be deployed in Layer 3 fabric, VxLAN overlay, or MC-LAG configurations.

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 OS; 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 process (daemon), alongside Junos OS. They can also communicate with Junos OS as well as 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, customers can implement specific and unique network forwarding functions and automate the control and management by integrating with their own SDN controller and orchestration tools. The QFX10000

programmable APIs have a stable abstraction so that the customer’s applications can work across Juniper platforms as well as future Junos OS releases. Based on Apache Thrift, the QFX10000 switches support standardized API definition via Interface Description Language (IDL), while allowing users to employ their favorite programming language.

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

Junos Space Network Director

The QFX10000 switches can be managed through Juniper Networks 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 and ensure faster rollout and activation of services.




Figure 3: Juniper virtualized open network operating system framework

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 L2, L3, and security services in the data center without the need for manual operator intervention.

Features and Benefits

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

QFX10000 switches support a number of network automation features, including operations and event scripts, automatic rollback, ZTP, and Python scripting. The switches also integrate 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, 6VPE), RSVP traffic engineering, and LDP to allow standards-based network segmentation and virtualization.


VXLAN

QFX10000 switches support L2 and L3 gateway services that enable VXLAN-to-VLAN connectivity at any tier of the data center network, from server access to the edge. The QFX10000 switches can integrate with VMware NSX through data plane (VXLAN) and control and management plane (OVSDDB) protocols to centrally automate and orchestrate the data center network.

FCoE

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

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



Specifications

Hardware

Table 1: QFX10002 System Capacity

	QFX10002-36Q	QFX10002-72Q
System throughput	Up to 2.88 Tbps	Up to 5.76 Tbps
Forwarding capacity	Up to 1 Bpps	Up to 2 Bpps
Maximum 10GbE port density	144	288
Maximum 40GbE port density	36	72
Maximum 100GbE port density	12	24

Table 2: QFX10002 System Specifications

	QFX10002-36Q	QFX10002-72Q
Dimensions (W x H x D):	17.4 x 3.46 x 31 in (44.2 x 8.8 x 78.7 cm)	17.4 x 3.46 x 31 in (44.2 x 8.8 x 78.7 cm)
Rack units	2 U	
Weight	50 lbs. (23 Kg)	68.6 lbs. (31 Kg)
Operating System	Junos OS	
CPU	Intel Quad Core Ivy Bridge 2.5 GHz CPU, 16 GB SDRAM	
Management and Precision Time Protocol (PTP) interfaces	<ul style="list-style-type: none"> 1 small form-factor pluggable transceiver (SFP/SFP+) port for PTP Grandmaster Fiber (SFP) or 10/100/1000BaseT (RJ45) Ethernet management port SMB In, SMB Out, 10 MHz In, 10 MHz Out One Console port USB 2.0 storage interface 	
Power	2x 1600 AC/DC power supply	4x 1600 AC/DC power supply
Cooling	<ul style="list-style-type: none"> Front-to-back airflow 3 hot-swappable fan modules with redundant fans 	
Total packet buffer	12 GB	24 GB
Latency	<ul style="list-style-type: none"> As low as 2.5 microseconds within a PFE As low as 5.5 microseconds across PFES 	
Warranty	Juniper standard one-year warranty	

Software

Table 3: QFX10002 Scale (Unidimensional)

	QFX10002-36Q	QFX10002-72Q
Media access control (MAC) addresses	Up to 256,000; up to 96,000 learned per PFE	Up to 512,000; up to 96,000 learned per PFE
ARP entries	Up to 192,000, or 64,000 per PFE	Up to 340,000, or 64,000 per PFE
Jumbo frames	9,216 bytes maximum	
VLANs	4,000	
Forwarding information base (FIB) (IPv4/IPv6)	Up to 2 million	
Routing Information Base (RIB) (v4)	10 million	
Routing Information Base (RIB) (v6)	4 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	
Filter terms	Up to 64,000	
Policers	Up to 8,000	
Output queues per port	8	
VOQs	384,000 per PFE	
LAGs	72	144
Members per LAG	64	
Generic routing encapsulation (GRE) tunnels	2,000	
MPLS L3 VPNs (Advanced Feature License)	4,000	

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)
- 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 RPF (uRPF)
- Loop-free alternate (LFA)
- BGP (Premium Feature license or Advanced Feature license)
- IS-IS (Premium Feature license or Advanced Feature license)
- Dynamic Host Configuration Protocol (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
- 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
 - 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 LSPs
- IPv4 Layer 3 VPNs
- IPv6 Layer 3 VPNs (6VPE)
- Layer 2 circuits

Overlays (Premium Feature License or Advanced Feature License)

- Integration with Juniper Networks Contrail*
- Integration with VMware NSX SDN controller*
- VXLAN

- Open vSwitch Database (OVSDDB)
- VXLAN L2 Gateway
- VXLAN L3 Gateway
- Ethernet VPN (EVPN)-VXLAN

High Availability

- GRES
- BFD

Timing

- Precision Timing Protocol (PTP)
 - Transparent Clock
 - Boundary Clock*

Visibility

- Switched Port Analyzer (SPAN)
- Remote SPAN (RSPAN)
- Encapsulated Remote SPAN (ERSPAN)
- Firewall filter-based port mirroring
- sFlow v5

Data Center Bridging

- Priority-based flow control (PFC)—IEEE 802.1Qbb
- 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)
- 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 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 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 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

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
- 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

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)

- QFX10002-36Q:
 - Maximum power draw: 800 W (AC, DC), 2,730 BTU/hr
 - Typical power draw: 560 W (AC, DC), 1,911 BTU/hr
- QFX10002-72Q:
 - Maximum power draw: 1,425 W (AC, DC), 4,862 BTU/hr
 - Typical power draw: 1,050 W (AC, DC), 3,583 BTU/hr

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/NZS 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 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

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

Product Number	Description
QFX10002 Hardware	
QFX10002-72Q	QFX10002 system with 72-port 40GbE QSFP+ / 24-port 100GbE QSFP28 / 288-port 10GbE SFP+ with 4 1600 W AC power supplies and 3 fan trays
QFX10002-36Q	QFX10002 system with 36-port 40GbE QSFP+ / 12-port 100GbE QSFP28 / 144-port 10GbE SFP+ with 2 1600 W AC power supplies and 3 fan trays
QFX10002-72Q-DC	QFX10002 system with 72-port 40GbE QSFP+ / 24-port 100GbE QSFP28 / 288-port 10GbE SFP+ with 4 1600 W DC power supplies and 3 fan trays
QFX10002-36Q-DC	QFX10002 system with 36-port 40GbE QSFP+ / 12-port 100GbE QSFP28 / 144-port 10GbE SFP+ with 2 1600 W DC power supplies and 3 fan trays
JPSU-1600W-DC-AFO	QFX10002 1600 W DC power supply
JPSU-1600W-AC-AFO	QFX10002 1600 W AC power supply
QFX10002-FAN	QFX10002 fan tray

QFX10002 Pluggable Optics

JNP-QSFP-100G-SR4	QSFP28 100GBASE-SR4 optics for up to 100 m transmission over parallel multimode fiber-optic (MMF)
JNP-QSFP-100G-LR4	QSFP28 100GBASE-LR4 optics for up to 10 km transmission over serial single-mode fiber-optic (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	100GbE 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 optic, 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
JNP-QSFP-DAC-10MA	QSFP+ to QSFP+ Ethernet direct attach copper (DAC) (twinax copper cable) 10 m active

Product Number	Description
JNP-QSFP-DAC-7MA	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 7 m active
JNP-QSFP-DAC-5M	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 5 m passive
JNP-QSFP-DAC-5MA	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 5 m active
QFX-QSFP-DAC-3M	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 3 m passive
QFX-QSFP-DAC-1M	QSFP+ to QSFP+ Ethernet DAC (twinax copper cable) 1 m passive
JNP-QSFP-DACBO-10M	QSFP+ to SFP+ 10-Gigabit Ethernet direct attach break out copper (twinax copper cable) 10 m active
JNP-QSFP-DACBO-7MA	QSFP+ to SFP+ 10-Gigabit Ethernet direct attach break out copper (twinax copper cable) 7 m active
JNP-QSFP-DACBO-5MA	QSFP+ to SFP+ 10-Gigabit Ethernet direct attach break out copper (twinax copper cable) 5 m active
QFX-QSFP-DACBO-3M	QSFP+ to SFP+ 10-Gigabit Ethernet direct attach break out copper (twinax copper cable) 3 m
QFX-QSFP-DACBO-1M	QSFP+ to SFP+ 10-Gigabit Ethernet direct attach break out copper (twinax copper cable) 1 m
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, 20 m on OM2, 100 m on OM3 multimode fiber
QFX-SFP-DAC-10MA	SFP+ 10-Gigabit Ethernet DAC (active twinax copper cable) 10 m
QFX-SFP-DAC-7MA	SFP+ 10-Gigabit Ethernet DAC (active twinax copper cable) 7 m

Product Number	Description
QFX-SFP-DAC-5MA	SFP+ 10-Gigabit Ethernet DAC (active twinax copper cable) 5 m
QFX-SFP-DAC-3M	SFP+ 10-Gigabit Ethernet DAC (twinax copper cable) 3 m
QFX-SFP-DAC-3MA	SFP+ 10-Gigabit Ethernet DAC (active twinax copper cable) 3 m
QFX-SFP-DAC-1M	SFP+ 10-Gigabit Ethernet DAC (twinax copper cable) 1 m
QFX-SFP-DAC-1MA	SFP+ 10-Gigabit Ethernet DAC (active twinax copper cable) 1 m
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

QFX10002 Software Feature Licenses

QFX10002-36Q-PFL	QFX10002-36Q Premium Feature License
QFX10002-36Q-AFL	QFX10002-36Q Advanced Feature License
QFX10002-72Q-PFL	QFX10002-72Q Premium Feature License
QFX10002-72Q-AFL	QFX10002-72Q Advanced Feature License

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