


# QFX5120 Ethernet Switch



## Product Overview

The QFX5120 delivers low-latency, native 25GbE, rich Layer 2 and Layer 3 features, VXLAN overlay support, and 100GbE uplinks, making it the ideal switch for access and top-of-rack deployments.

Featuring Layer 3 gateway capabilities for bridging between virtualized and bare-metal servers, the QFX5120 is designed for extremely agile data centers that demand support for overlay/underlay network architectures. The high-density 10GbE, 25GbE, 40GbE, and 100GbE ports also make the QFX5120 ideally suited for deployment as leaf-and-spine and leaf topologies.

## Product Description

Data centers are expanding rapidly as they adopt more cloud services, creating a need for higher network bandwidth per rack. As a result, 25GbE and 100GbE have become the dominant port speeds in most data centers. The tremendous growth of cloud services, coupled with the widespread adoption of overlay technologies, has created a need for highly agile switching platforms that can satisfy the demands of these evolving data centers.

The high-performance Juniper Networks® QFX5120 Ethernet Switch fits the bill, providing the foundation for dynamic data centers. Native 25GbE ports and 100GbE uplinks satisfy growing demand for 100GbE spine ports to support all server access speeds, including 10GbE and 25GbE. The QFX5120 also supports a diverse set of switching architectures, including data center fabric, Layer 3, and spine-and-leaf deployments, enabling users to easily adapt and evolve as requirements change over time.

## Architecture and Key Components

The QFX5120 includes native 25GbE (fiber only) and 40GbE or 100GbE fixed-configuration options with rich L2, L3, Ethernet VPN (EVPN), and MPLS features. The QFX5120 runs the same reliable, high-performance Juniper Networks Junos® operating system used by all Juniper Networks QFX Series Switches and EX Series Ethernet Switches, SRX Series Services Gateways, and Juniper Networks routers, ensuring a consistent control plane feature implementation and user experience across the entire Juniper infrastructure.

## Product Options

The QFX5120 switch is a compact, 1 U platform that provides wire-speed packet performance, very low latency, and a rich set of Junos OS features. In addition to a high-throughput Packet Forwarding Engine (PFE), the performance of the QFX5120 control plane is enhanced with a powerful 2.2 GHz quad-core Intel CPU with 16 GB of memory and 64 GB SSD storage.

As the first member of the QFX5120 line, the QFX5120-48Y is a 25GbE/100GbE data center access switch offering 48 small form-factor pluggable plus (SFP+) transceiver ports and eight quad SFP (QSFP28) ports that can be configured as 8x40GbE or 8x100GbE ports, with an aggregate throughput of 2 Tbps or 2 Bpps per switch. For added flexibility, each QSFP28 port can also be configured as 4x25GbE ports using breakout cables, while each QSFP+ port can be configured as 4x10GbE ports using breakout cables, increasing the number of supported 25GbE and 10GbE ports per switch to 80.

## QFX5120-48Y Highlights

The QFX5120-48Y switch features the following:

- Support for high-density, multispeed 10/25/40/100GbE access and aggregation deployments, with up to 48 25GbE ports and eight 100GbE uplink ports in a 1 U platform
- Up to 2 Tbps L2 and L3 performance, with latency as low as 550 nanoseconds
- A 2.2 GHz quad-core Intel CPU with 16 GB memory and 64 GB SSD storage
- Feature-rich automation capabilities with support for Python and Zero Touch Provisioning (ZTP)
- Support for Virtual Extensible LAN (VXLAN) as an L2 or L3 gateway and Open vSwitch Database (OVSDb) protocol as an L2 gateway
- Advanced Junos OS features such as EVPN, BGP add-path, MPLS, L3 VPN, and IPv6 6PE

## Junos OS

The high-performance QFX5120 switch runs Junos OS, Juniper's powerful and robust network operating system that powers all Juniper switches, routers, and firewalls. Key Junos OS features that enhance the functionality and capabilities of the QFX5120 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 active routing (NSR) and nonstop bridging (NSB)
- Commit and rollback functionality that ensures error-free network configurations
- A powerful set of scripts for on-box problem detection, reporting, and resolution

## Junos OS Software License

The software features supported on the QFX5120-48Y switch are categorized into four tiers: Base, Premium-Lite, Premium, and Advanced.

- Base software features include basic L2 switching, basic L3 routing, multicast, automation, programmability, ZTP, and basic monitoring. A Base software feature license comes with the purchase of the hardware and does not require any explicit license keys.
- Premium-Lite software features include BGP and IS-IS to explicitly address the needs of IP fabric deployments at cloud service providers. To enable these features, customers must purchase the QFX5K-C1-PFL-LITE license, generate unique license keys, and install them on the switch. The license is not portable across devices.
- Premium software features include BGP, IS-IS, and EVPN VXLAN to explicitly address the needs of enterprise customers. To enable these features, customers must

purchase the QFX5K-C1-PFL license, generate unique license keys, and install them on the switch. The license is not portable across devices.

- Advanced software features include all Premium license functionality plus MPLS to explicitly address the needs of data center interconnect and edge use cases. To enable these features, customers must purchase the QFX5K-C1-AFL license, generate unique license keys, and install them on the switch. The license is not portable across devices.

The Premium and Advanced software licenses are classified as Class 1, Class 2, Class 3, etc., and are offered as perpetual licenses. Class 1 licenses are applicable to the QFX5120-48Y switch. Please see the Ordering Information section for license SKU descriptions.

## Data Center Deployments

Data centers demand high-speed, low-latency, storage-converged and I/O-converged networking solutions that can maximize performance for physical servers, virtual servers, and storage. The QFX5120 switch addresses these issues with low-latency, lossless, high-density 10GbE and 25GbE interfaces, as well as 100GbE uplinks to the core network. Furthermore, the QFX5120 offers EVPN VXLAN Layer 2 and Layer 3 gateway support, making it an ideal solution for overlay deployments in the data center. The QFX5120 also supports flexible back-to-front and front-to-back airflow cooling options, ensuring consistency with server designs for hot-aisle or cold-aisle deployments.

## Data Center Server Access

The QFX5120 switch supports tri-speed 1/10/25GbE, making it a perfect fit for top-of-rack deployments. The 48 ports of native 10/25GbE for server connectivity, plus up to eight 40GbE or 100GbE ports for uplink connectivity, provide very low oversubscription of 1.5:1 from access to aggregation. Each 40GbE port can also be broken out into four 10GbE ports, providing additional options for server connectivity.

The QFX5120 can operate in both cut-through and store-and-forward modes, delivering sustained wire-speed switching with sub-microsecond latency and low jitter for any packet size (including jumbo frames) in either mode. With features such as multichassis link aggregation group (MC-LAG), the QFX5120 supports active/active server dual homing and can utilize full bisectonal bandwidth from server to switch.

Equipped with Junos OS, the QFX5120 supports the most advanced and robust routing capabilities in the industry, including RIP and OSPF for both IPv4 and IPv6 as well as advanced routing capabilities such as IS-IS and BGP. With additional capabilities such as 64-way equal-cost multipath (ECMP) and BGP add path, the QFX5120 is an ideal building block for deploying the most robust Layer 3 underlay for SDN.

## Features and Benefits

### Automation

The QFX5120 supports a number of network automation and plug-and-play operation features, including ZTP and event scripts, automatic rollback, and Python scripting.

### Junos Telemetry Interface

The QFX5120 supports Junos Telemetry Interface (JTI), a modern telemetry streaming tool designed for performance monitoring in complex, dynamic data centers. Streaming data to a performance management system enables network administrators to measure trends in link and node utilization, and troubleshoot such issues as network congestion in real time. JTI delivers the following features:

- Application visibility and performance management by provisioning sensors to collect and stream data and analyze application and workload flow path through the network
- Capacity planning and optimization by proactively detecting hotspots and monitoring latency and microbursts
- Troubleshooting and root cause analysis via high-frequency monitoring and correlation of overlay and underlay networks

### Monitoring and Analytics with Juniper AppFormix

Juniper AppFormix®, a network and device monitoring platform, provides comprehensive visualization and smart analytics for the QFX5120 switch. Powered by machine learning and artificial intelligence, AppFormix automatically learns network and device performance profiles, detects faults, and facilitates preventative remediation in real time at scale.

Based on a distributed policy and analytics architecture, AppFormix seamlessly collects and analyzes telemetry in real time from the QFX5120 platform via JTI, generating events and alert notifications and providing the operator with actionable insights.

AppFormix delivers the following features and benefits:

- **Real-time metrics and alarms:** AppFormix collects and analyzes comprehensive metrics in real time for monitored networks and devices, displaying them in the AppFormix dashboard as real-time charts and graphs. With AppFormix, an operator can also configure alarms for individual or groups of metrics that are triggered when specific values or conditions are met, or when anomalies occur, notifying the operator in real time. AppFormix software also plots these alarms on both real-time and historical charts, giving the operator a holistic view of current status as well as prevailing trends.
- **SLA Monitor:** The AppFormix dashboard includes an SLA Monitor pane that provides operators with an overview showing the current state of the network and devices. The SLA status indicates whether the network and the device are currently operating outside a user-defined performance threshold, while risk predicts whether the device may be unhealthy in the future based on historical trends.
- **Notifications:** AppFormix can be configured to notify external systems when events and alarms are triggered. For each event or alarm, AppFormix can post a structured description of the event as a JavaScript Object Notation (JSON) payload to an external HTTP endpoint. These

## AppFormix Operations tools for the DevOps era

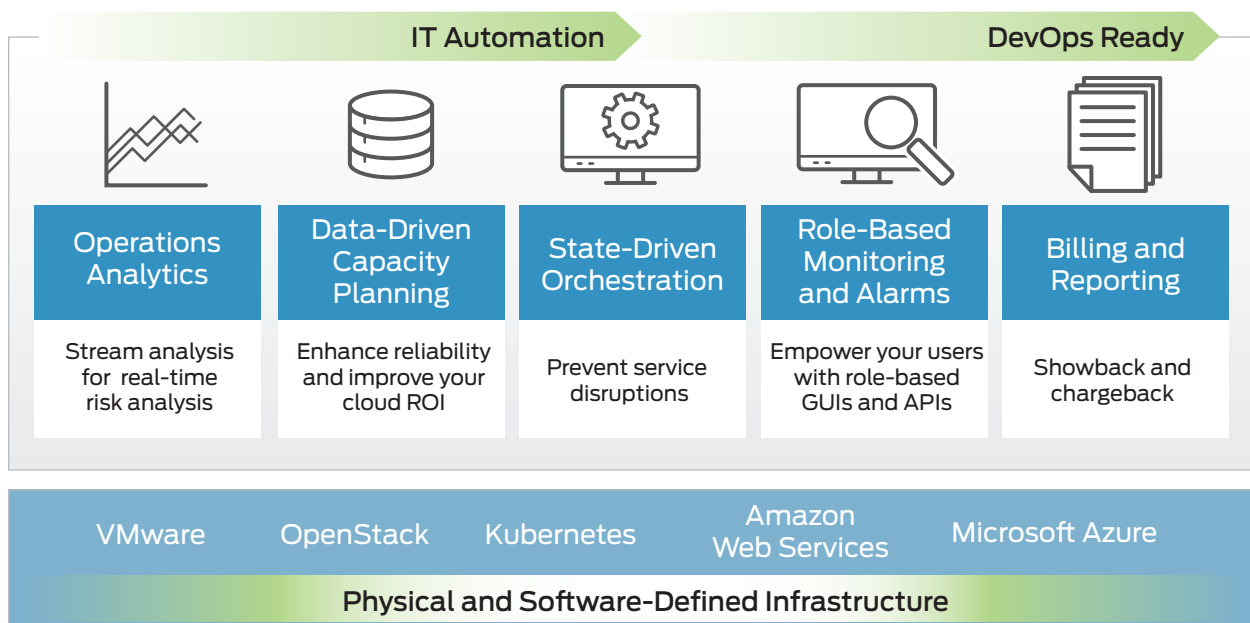


Figure 1. AppFormix overview

notifications can initiate an action or workflow, whether corrective, preventive, or otherwise, to keep the network and devices operating in an optimal state. For operator convenience, AppFormix comes preconfigured with some of the industry’s most popular incident management and collaboration platforms, including PagerDuty, ServiceNow, and Slack.

- **Reports:** AppFormix reports allow users to analyze how networks and devices are being consumed over time. The dashboard displays reports in both graphical and tabular formats. Users may also download the data as an HTML-formatted report, raw comma-separated values (CSV) file, or JSON-formatted statistics for further analysis.
- **Flexible forwarding table:** The QFX5120 switch includes a Unified Forwarding Table (UFT), which 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 L2 environment, the QFX5120 supports 288,000 MAC addresses. In L3 mode, the table can support 208,000 host entries. And in LPM mode, it can support 351,000 prefixes. Junos OS provides configurable options through a CLI that can optimize the QFX5120 for various deployment scenarios.
- **Intelligent buffer management:** The QFX5120 features a total of 32 MB of shared buffers. While 25% of the total buffer space is dedicated, the rest is shared among all ports and is user-configurable. The intelligent buffer mechanism in the QFX5120 effectively absorbs traffic bursts while providing deterministic performance, significantly increasing performance over static allocation.
- **MPLS:** A broad set of MPLS features, including L3 VPN, IPv6 provider edge router (6PE), RSVP traffic engineering, and LDP allow standards-based network segmentation and virtualization, enabling the QFX5120 to be deployed as a low-latency MPLS label-switching router (LSR).
- **VXLAN overlays:** The QFX5120 switch is capable of both L2 and L3 gateway services. Customers can deploy overlay networks to provide L2 adjacencies for applications over L3 fabrics. The overlay networks use VXLAN in the data plane and EVPN or OVSDB for programming the overlays, which can operate without a controller or be orchestrated with an SDN controller like OpenContrail.

### QFX5120 Deployment Options

The following table shows just two of the many QFX5120 deployment possibilities.

Port Combinations	Switch	Deployment
48x10GbE + 8x40GbE	QFX5120-48Y	10GbE Access or Leaf
48x25GbE + 8x100GbE	QFX5120-48Y	25GbE Access or Leaf

The following figures show various data center architectures where the QFX5120 serves as a network building block. In Figure 2, the QFX5120 is deployed as a leaf that is acting as an edge-routed gateway. In this topology, the VXLAN tunnel encapsulation and decapsulation take place on the QFX5120 leaf switches, while QFX5200-32C or QFX5210-64C spine switches are used as an IP transit. QFX5110-32Q switches can also be used in the spine to build a 40GbE fabric.

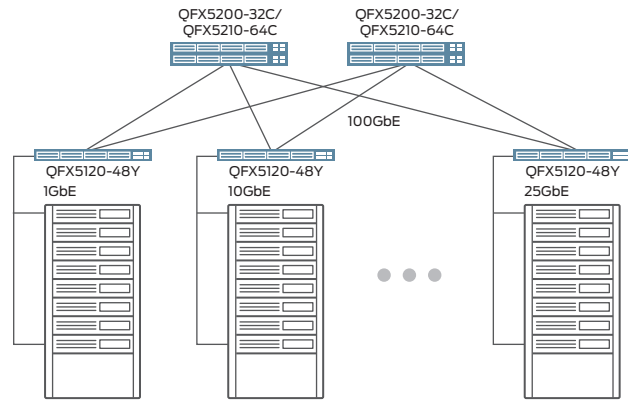


Figure 2. Edge-routed gateway at leaf with QFX5200/QFX5210 as spine

In Figure 3, the QFX5120 is deployed as a top-of-rack switch with QFX10000 spine switches acting as centrally routed gateways. In this topology, the VXLAN tunnel encapsulation and decapsulation occur on the QFX10000 spine switches. In addition to the QFX5120, QFX5110 and QFX5200-32C switches can also be configured as leaf nodes in this deployment.

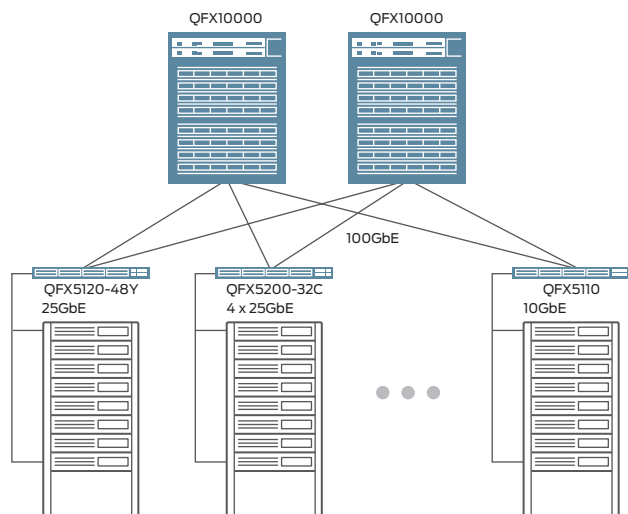



Figure 3. QFX5120-48Y in a leaf-and-spine deployment



## QFX5120 Switch Specifications

### Hardware

#### Switching Capacity

- 2 Tbps/1.31 Bpps

#### Weight

- 23.7 lb (10.75 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 (airflow out) for hot aisle deployment
- Back-to-front (airflow in) for cold aisle deployment
- Management and console port connections

#### Power Consumption

- Max load: 450 W
- Typical load: 260 W
- Idle load: 160 W

#### Interface Options

- 1GbE SFP: 48 (24 copper 1GbE)
- 10GbE SFP+: 48/80 (with breakout cable)
- 25GbE SFP: 48/80 (with breakout cable)
- 40GbE QSFP+: 8
  - Each QSFP+ port can be configured as a 4 x 10GbE interface or as a 40 Gbps port
- 100GbE QSFP28: 8
  - Each QSFP28 port can be configured as a 4 x 25GbE interface or as a 100 Gbps port
- 1 USB 2.0 port
- 1 RS-232 console port
- 2 management ports: 2x RJ-45 ports
- Supported transceiver and direct attach cable
- SFP GbE optical and copper module
- SFP+ 10GbE optical modules
- SFP+ DAC cables: 1/3/5 m twinax copper and 1/3/5/7 m active twinax copper
- SFP28 DAC cables: 1/3 m twinax copper
- SFP28 Optics: SR, LR
- QSFP+ DAC cables: 1/3 m twinax copper
- QSFP+ Optics: SR4, LX4, ESR4, ER4, LR4
- QSFP+ to SFP+: 10GbE direct attach breakout copper (1/3 m twinax copper cable)
- QSFP28 Optics: SR4, ER4, PSM4, CWDM4, LR4

#### Rack Installation Kit

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

#### Airflow

- Redundant (N+1) and hot-pluggable fan modules for front-to-back and back-to-front airflow
- Redundant variable-speed fans to reduce power draw

### Power Supply and Fan Modules

- Dual redundant (1+1) and hot-pluggable 650 W AC/DC power supplies
- 100-240 V single phase AC power
- -48 to -60 V DC power supply
- Redundant (4+1) and hot-pluggable fan modules for front-to-back or back-to-front airflow

### Performance Scale (Unidimensional)

- MAC addresses per system: 288,000
- VLAN IDs: 4093
- Number of link aggregation groups (LAGs): 80
- Number of ports per LAG: 64
- IPv4 unicast routes: 351,000 prefixes; 208,000 host routes; 64 ECMP paths
- IPv4 multicast routes: 104,000
- IPv6 unicast routes: 168,000 prefixes; 104,000 host routes
- IPv6 multicast routes: 52,000
- ARP entries: 64,000
- Jumbo frame: 9216 bytes
- Spanning Tree Protocol (STP)
  - Multiple Spanning Tree Protocol (MSTP) Instances: 64
  - VLAN Spanning Tree Protocol (VSTP) Instances: 509
- Traffic mirroring
  - Mirroring destination ports per switch: 4
  - Maximum number of mirroring sessions: 4
  - Mirroring destination VLANs per switch: 4

### Software Features Supported

#### 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
- Private VLAN (PVLAN)
- VLAN translation
- 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)



### Link Aggregation

- Multichassis link aggregation (MC-LAG)
- LAG load sharing algorithm—bridged or routed (unicast or multicast) traffic:
  - IP: SIP, Dynamic Internet Protocol (DIP), TCP/UDP source port, TCP/UDP destination port
  - Layer 2 and non-IP: MAC SA, MAC DA, Ethertype, VLAN ID, source port

### Layer 3 Features (IPv4)

- Static routing
- 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, and v3 (Layer 2 only)
- IGMP Filter
- PIM-SM, PIM-SSM, PIM-DM
- 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
- Sticky MAC address
- DHCP snooping

### Quality of Service (QoS)

- L2 and L3 QoS: Classification, rewrite, queuing
- Rate limiting:
  - Ingress policing: Single-rate two-color policer, two-rate three-color policer
  - Egress policing: Policer, policer mark down action
  - Egress shaping: Per queue on each port
- 10 hardware queues per port (8 unicast and 2 multicast)
- Strict-priority queuing (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 (ingress)
- Remarking of bridged packets

### Data Center Bridging (DCB)

- Priority-based flow control (PFC)—IEEE 802.1Qbb
- Data Center Bridging Capability Exchange (DCBX), DCBx Fibre Channel over Ethernet (FCoE), and iSCSI type, length, and value (TLVs)

### High Availability

- Bidirectional Forwarding Detection (BFD)
- Uplink failure detection

### 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 LSR support
- IPv6 tunneling (6PE) (via IPv4 MPLS backbone)
- IPv4 L3 VPN (RFC 2547, RFC 4364)

### Server Virtualization Management and SDN-Related Protocols

- VXLAN OVSDB
- EVPN VXLAN

### Management and Operations

- 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
- Beacon LED for port and system
- ZTP
- OpenStack Neutron Plug-in
- Python
- Junos OS event, commit, and OP scripts
- Junos Telemetry Interface (JTI)

### Traffic Mirroring

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

## Standards Compliance

### IEEE Standard

- IEEE standard
- IEEE 802.1D
- IEEE 802.1w
- IEEE 802.1
- IEEE 802.1Q
- IEEE 802.1p
- IEEE 802.1ad
- IEEE 802.3ad
- IEEE 802.1AB
- IEEE 802.3x
- IEEE 802.1Qbb
- IEEE 802.1Qaz

### T11 Standards

- INCITS T11 FC-BB-5

### Supported RFCs

- RFC 768 UDP
- RFC 783 Trivial File Transfer Protocol (TFTP)
- RFC 791 IP
- RFC 792 ICMP
- RFC 793 TCP
- RFC 826 ARP
- RFC 854 Telnet client and server
- RFC 894 IP over Ethernet
- RFC 903 RARP
- RFC 906 TFTP Bootstrap
- RFC 951 1542 BootP
- RFC 1058 Routing Information Protocol
- RFC 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 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 with 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 LSA Option
- RFC 2385 Protection of BGP Sessions via the TCP MD5 Signature Option
- RFC 2439 BGP Route Flap Damping
- RFC 2453 RIP v2
- RFC 2474 Definition of the Differentiated Services Field (DS 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 the Transmission Control Protocol using SMIv2
- RFC 2013 SNMPv2 for the 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) Applications—(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. 62368-1-14 Information Technology Equipment - Safety
- UL 62368-1 Information Technology Equipment – Safety
- EN 62368-1: 2014 Information Technology Equipment – Safety
- IEC 62368-1: 2014 2nd Edition Information Technology Equipment - Safety (All country deviations): CB Scheme

### EMC

- EN 300 386 V1.6.1 (2012-09) Electromagnetic compatibility and Radio spectrum Matters (ERM) Telecommunication network equipment
- EN 300 386 V2.1.1 (2016-07) Telecommunication network equipment; EMC requirements; Harmonized Standard covering the essential requirements of the Directive 2014/30/EU
- EN 55032:2012 (CISPR 32:2012) Electromagnetic compatibility of multimedia equipment—Emission requirements
- EN 55024:2010 (CISPR 24:2010) Information technology equipment—Immunity characteristics—Limits and methods of measurement
- IEC/EN 61000 Immunity Test
- AS/NZS CISPR 32:2015 Australia/New Zealand Radiated and Conducted Emissions
- FCC 47 CFR Part 15 USA Radiated and Conducted Emissions
- ICES-003 Canada Radiated and Conducted Emissions
- VCCI-CISPR 32:2016 Japanese Radiated and Conducted Emissions
- BSMI CNS 13438 Taiwan Radiated and Conducted Emissions (at 10 meters)
- KN32/KN35 Korea Radiated Emission and Immunity Characteristics (at 10 meters)
- KN61000 Korea Immunity Test
- TEC/SD/DD/EMC-221/05/OCT-16 India EMC standard

## 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 6000 ft (1829 m)
- Relative humidity operating: 5% to 90% (noncondensing)
- Relative humidity nonoperating: 0% to 95% (noncondensing)

## Juniper Networks Services and Support

Juniper Networks leads the market in performance-enabling services designed to accelerate, extend, and optimize your deployments. Our services enable you to maximize operational efficiency, reduce costs, and minimize risk while achieving a faster time-to-value for your network.

By leveraging best practices from across the industry, you get the maximum levels of system performance, designed and delivered by the world's leading professional technology experts.

For more information, please visit [www.juniper.net/us/en/products-services](http://www.juniper.net/us/en/products-services).

## Installation and Implementation Service

Juniper Professional Services offers a Data Center Switching QuickStart program to ensure that the solution is operational and the customer has a complete understanding of areas such as configuration and ongoing operations. The QuickStart service provides an onsite consultant who works with the client team to quickly develop the initial configuration and deployment of a small Juniper Networks data center switching environment. A knowledge transfer session, which is intended as a review of local implementation and configuration options, is also included, but is not intended as a substitute for formalized training.

## Ordering Information

Product	Description
<b>Switch Hardware</b>	
QFX5120-48Y-AFI	QFX5120, 48x25GbE+8x100GbE 1 U AC airflow in
QFX5120-48Y-AFO	QFX5120, 48x25GbE+8x100GbE 1 U AC airflow out
QFX5120-48Y-DC-AFI	QFX5120, 48x25GbE+8x100GbE 1 U DC airflow in
QFX5120-48Y-DC-AFO	QFX5120, 48x25GbE+8x100GbE 1 U DC airflow out
JPSU-650W-AC-AO	Juniper 650 W AC power supply (port-side-to-FRU-side airflow)
JPSU-650W-AC-AI	Juniper 650 W AC power supply (FRU-side-to-port-side airflow)
JPSU-650W-DC-AFO	Juniper 650 W DC power supply (port-side-to-FRU-side airflow)
JPSU-650W-DC-AFI	Juniper 650 W DC power supply (FRU-side-to-port-side airflow)
QFX5110-FANAFI	FANAFI fan model, back-to-front airflow

Product	Description
QFX5110-FANAFO	FANAFO fan model, front-to-back airflow
EX-4PST-RMK	4 post rack mount

## Optics and Transceivers

QFX-SFP-1GE-T	SFP 1000BASE-T copper transceiver module for up to 100 m transmission on Category 5
QFX-SFP-1GE-SX	SFP 1000BASE-SX GbE optics, 850 nm for up to 550 m transmission on multimode fiber
QFX-SFP-1GE-LX	SFP 1000BASE-LX GbE optics, 1310 nm for 10 km transmission on single-mode fiber
QFX-SFP-10GE-USR	SFP+ 10GbE ultra short reach optics, 850 nm for 10 m on OM1, 20 m on OM2, 100 m on OM3 multimode fiber
QFX-SFP-10GE-SR	SFP+ 10GBASE-SR 10GbE optics, 850 nm for up to 300 m transmission on multimode fiber
QFX-SFP-10GE-LR	SFP+ 10GBASE-LR 10GbE optics, 1310 nm for 10 km transmission on single-mode fiber
QFX-SFP-10GE-ER	SFP+ 10GBASE-ER 10GbE optics, 1550 nm for 40 km transmission on single-mode fiber
EX-SFP-10GE-ZR	SFP+ 10GBASE-ZR 10GbE optics, 1550 nm for 80 km transmission on single-mode fiber
QFX-SFP-DAC-1M	SFP+ 10GbE Direct Attach Copper (twinax copper cable), 1 m
QFX-SFP-DAC-3M	SFP+ 10GbE Direct Attach Copper (twinax copper cable), 3 m
QFX-SFP-DAC-5M	SFP+ 10GbE Direct Attach Copper (twinax copper cable), 5 m
QFX-SFP-DAC-1MA	SFP+ 10GbE Direct Attach Copper (active twinax copper cable), 1 m
QFX-SFP-DAC-3MA	SFP+ 10GbE Direct Attach Copper (active twinax copper cable), 3 m
QFX-SFP-DAC-5MA	SFP+ 10GbE Direct Attach Copper (active twinax copper cable), 5 m
QFX-SFP-DAC-7MA	SFP+ 10GbE Direct Attach Copper (active twinax copper cable), 7 m
JNP-25G-SR	SFP28 25GBASE-SR Optics for up to 100 m transmission over serial multimode fiber-optic (MMF) OM4 fiber
JNP-25G-LR	SFP28 25GBASE-SR Optics for up to 10 km transmission over serial single-mode fiber-optic (SMF)
JNP-25G-DAC-1M	25GbE SFP to SFP copper cable, 1 m
JNP-25G-DAC-3M	25GbE SFP to SFP copper cable, 3 m
JNP-QSFP-40G-LX4	40GbE QSFP+ LX4 optics
QFX-QSFP-40G-SR4	QSFP+ 40GBASE-SR4 40GbE optics, 850 nm for up to 150 m transmission on multimode fiber
QFX-QSFP-40G-ESR4	QSFP+ 40GBASE-SR4 40GbE optics, 850 nm for up to 300 m transmission on multimode fiber
JNP-QSFP-40GE-ER4	QSFP+ 40GBASE-ER4 40GbE optics, 1310 nm for up to 40 km transmission on SMF
QSFP-40GBASE-ER4	40GBASE-ER4 QSFP+ pluggable transceiver
JNP-QSFP-40G-LR4	40GbE QSFP+ LR4

Product	Description
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
QFX-QSFP-DACBO-1M	QSFP+ to SFP+ 10GbE Direct Attach Breakout Copper (twinax copper cable), 1 m
QFX-QSFP-DACBO-3M	QSFP+ to SFP+ 10GbE Direct Attach Breakout Copper (twinax copper cable), 3 m
JNP-QSFP-100G-SR4	QSFP28 100GbE, SR4, 100 m
QSFP-100GBASE-SR4	100GBASE-SR4 QSFP28 pluggable module, support only Ethernet rate
JNP-QSFP-100G-PSM4	QSFP28 100GBASE-PSM4 optics for up to 500 m transmission over parallel SMF
JNP-QSFP-100G-CWDM	QSFP28 100GbE, CWDM4, 2 km
JNP-QSFP-100G-LR4	QSFP28 100GbE, LR4, 10 km
JNP-100G-DAC-1M	QSFP28 to QSFP28 Ethernet Direct Attach Copper (twinax copper cable), 1 m
JNP-100G-DAC-3M	QSFP28 to QSFP28 Ethernet Direct Attach Copper (twinax copper cable), 3 m
JNP-100G-4X25G-1M	QSFP28 to SFP+ 25GbE Direct Attach Breakout Copper (twinax copper cable), 1 m
JNP-100G-4X25G-3M	QSFP28 to SFP+ 25GbE Direct Attach Breakout Copper (twinax copper cable), 3 m

### Software Feature Licenses

QFX5K-C1-PFL	QFX5000 Class 1 Premium Feature License
QFX5K-C1-AFL	QFX5000 Class 1 Advanced Feature License
QFX5K-C1-PFL-LITE	QFX5000 Class 1 Premium-Lite Feature License

## About Juniper Networks

Juniper Networks brings simplicity to networking with products, solutions and services that connect the world. Through engineering innovation, we remove the constraints and complexities of networking in the cloud era to solve the toughest challenges our customers and partners face daily. At Juniper Networks, we believe that the network is a resource for sharing knowledge and human advancement that changes the world. We are committed to imagining groundbreaking ways to deliver automated, scalable and secure networks to move at the speed of business.

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](http://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 2018 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Juniper, and Junos 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