



# QFX5120 ETHERNET SWITCH DATASHEET

## Product Overview

The [QFX5120 Switch](#) delivers rich, low latency Layer 2/Layer 3 features and advanced [EVPN-VXLAN](#) capabilities, making it an ideal data center top-of-rack and distribution switch for campus enterprise deployments.

Featuring L3 gateway capabilities for routing between virtualized and bare-metal servers, the QFX5120 is designed for extremely agile [data centers](#) that require support for overlay/underlay network architectures. Native 25GbE with 100GbE uplink ports on the QFX5120-48Y and QFX5120-48YM, 10GbE/1GbE copper with 100GbE uplink ports on the QFX5120-48T, and 32 100GbE ports on the QFX5120-32C make the QFX5120 family ideal for spine-and-leaf network deployments.

## Product Description

The Juniper Networks® QFX5120 Switch delivers high scale, high availability, and high performance for data center and campus deployments. The QFX5120 Switch is a versatile [routing](#) and [switching](#) platform addressing higher server access speed and campus distribution use cases while offering high-density 1GbE/10GbE/25GbE and 100GbE uplinks for collapsed spine data center or campus distribution deployments.

## Product Options

The QFX5120 Switch includes four compact 1 U platforms—the QFX5120-48Y, the QFX5120-48YM, the QFX5120-48T, and the QFX5120-32C—that provide wire-speed packet performance, very low latency, and a rich set of [Junos® operating system](#) features.

### QFX5120-48Y

The QFX5120-48Y is a 25GbE/100GbE data center leaf and campus distribution switch featuring:

- 48 25GbE (SFP28)/10GbE (SFP+)/1GbE (SFP) downlink ports
- Eight 100GbE (QSFP28)/40GbE (QSFP+) uplink ports
- Up to 2/4 Tbps (uni/bidirectional) of L2 and L3 performance, with latency as low as 800 nanoseconds
- A 2.2 GHz quad-core Intel CPU with 16 GB memory and 50 GB SSD storage

Using breakout cables, each of the eight 100GbE QSFP28 ports can be broken into four 25GbE SFP28 ports, while each 40GbE quad small form-factor pluggable plus (QSFP+) transceiver ports can be broken into four 10GbE small form-factor pluggable plus (SFP+) transceiver ports, increasing the total number of supported 25GbE and 10GbE ports per switch to 80.

### QFX5120-48YM

The QFX5120-48YM is a 10GbE/25GbE/100GbE data center leaf and campus distribution switch featuring:

- 48 25GbE (SFP28)/10GbE (SFP+)/1GbE (SFP) downlink ports
- Eight 100GbE (QSFP28)/40GbE (QSFP+) uplink ports
- Up to 2/4 Tbps (uni/bidirectional) of L2 and L3 performance, with latency as low as 800 nanoseconds
- Media Access Control Security (MACsec) AES-256 support across all ports
- A 2.9 GHz quad-core Intel CPU with 16 GB memory and 100 GB SSD storage

When using breakout cables, two of the 100GbE QSFP28 ports (ports 50 and 52) can be broken into four 25GbE SFP28 or four 10GbE SFP+ ports, increasing the maximum number of 10GbE/25GbE ports supported to 56. When using breakouts, the total number of all ports on the switch, including 6x100GbE ports and 56x10/25GbE, is 62.

### QFX5120-48T

The QFX5120-48T is a 10GbE/100GbE data center leaf and campus distribution switch featuring:

- 48 dual-speed 1GbE/10GbE RJ-45 copper downlink ports
- Six dual-speed 100GbE (QSFP28)/40GbE (QSFP+) uplink ports
- Up to 1.08/2.16 Tbps (uni/bidirectional) of L2 and L3 performance, with latency as low as 800 nanoseconds
- A 2.2 GHz quad-core Intel CPU with 16 GB memory and 100 GB SSD storage

Using breakout cables, 40GbE QSFP+ ports 50 and 51 can be channelized into four 10GbE SFP+ ports each, increasing the total number of 10GbE ports per switch to 56.

### QFX5120-32C

The QFX5120-32C is a compact 100GbE data center leaf-and-spine and campus distribution switch featuring:

- 32 100GbE (QSFP28) or 40GbE (QSFP+) uplink ports
- A 2.2 GHz quad-core Intel CPU with 16 GB memory and 64 GB SSD storage
- Up to 3.2/6.4 Tbps (uni/bidirectional) L2 and L3 performance, with latency as low as 800 nanoseconds

Using breakout cables, 100GbE QSFP28 ports 0-30 can be channelized into four 25GbE SFP28 ports\*, and the 40GbE QSFP+ ports 0-30 can be channelized into four 10GbE SFP+ ports, increasing the total number of 25GbE per switch to 124 and the total number of 10GbE ports per switch to 126 (including SFP+ ports 32 and 33). Port 31 does not support 4x10G or 4x25G channelization options. Only 100GbE, 40GbE, and 2x50GbE channelization options are supported on port 31.

In addition, all QFX5120 switch models include:

- Support for VXLAN as an L2 or L3 gateway
- Advanced Junos OS features such as Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN), BGP add-path, L3 VPN, and MPLS
- Feature-rich automation capabilities with support for Python and zero-touch provisioning (ZTP)

### 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

### Data Center Deployments

Data centers demand high-speed, low-latency, storage- and I/O-converged networking solutions that maximize performance for physical servers, virtual servers, and storage. The QFX5120 Switch addresses these issues with low-latency, lossless, high-density 10GbE, 25GbE, and 100GbE interfaces on a compact 1 U platform. In addition, the QFX5120 offers EVPN-VXLAN L2 and L3 gateway support, making it an ideal solution for either edge routed or centrally routed 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 QFX5120-48Y and QFX5120-48YM support tri-speed 1GbE/10GbE/25GbE, making them a perfect fit for top-of-rack deployments. The 48 ports of native 10GbE/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.

The QFX5120-48T supports dual-speed 1GbE/10GbE, also making it a perfect fit for top-of-rack deployments. The 48 native 10GbE RJ-45 copper ports for server connectivity, plus up to six 40GbE or 100GbE ports for uplink connectivity, provide an unsubscribed (0.8:1) access-to-aggregation ratio.

The QFX5120-32C can also be used for high-density 25GbE server connectivity, with an option to break out the 100GbE ports into four 25GbE ports.

In Figure 1, the QFX5120 is deployed as a leaf acting as an edge-routed gateway. In this topology, the VXLAN tunnel encapsulation and decapsulation take place on the QFX5120 leaf switches, while Juniper Networks [QFX5200-32C](#) or Juniper Networks [QFX5210-64C spine switches](#) are used for IP transit. Juniper Networks [QFX5110-32Q switches](#) can also be used in the spine to build a 40GbE fabric.

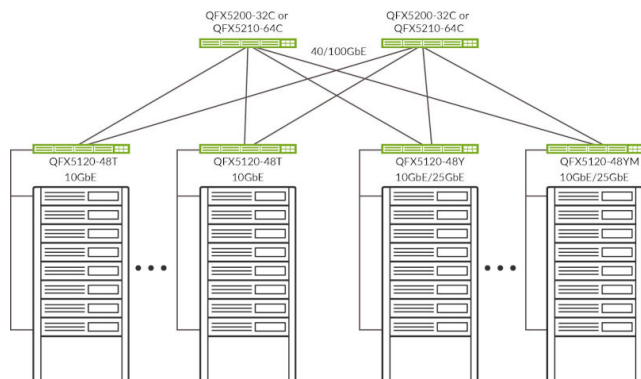


Figure 1: Edge-routed bridging at the leaf with QFX5120 as leaf switches and QFX5200-32C/QFX5210-64C as spine switches

In Figure 2, the QFX5120-48YM leaf and [Juniper Networks QFX10008 Switch](#) spine are deployed as EVPN-VXLAN switches acting as centrally routed gateways or distributed edge-routed gateways. If centrally routed bridging is used, the VXLAN tunnel encapsulation and decapsulation occur on the spine switches for inter-IRB (integrated routing and bridging) symmetric routing purposes. If edge-routed bridging is used, the IP first hop gateways are distributed at the leaf-level QFX5120-48YM switches using Type 5 symmetric inter-IRB routing. When using a QFX10000-30C-M MACsec line card installed in a spine QFX10008 Switch working with a leaf QFX5120-48YM switch, the leaf-to-spine architecture offers end-to-end MACsec AES-256 capabilities.

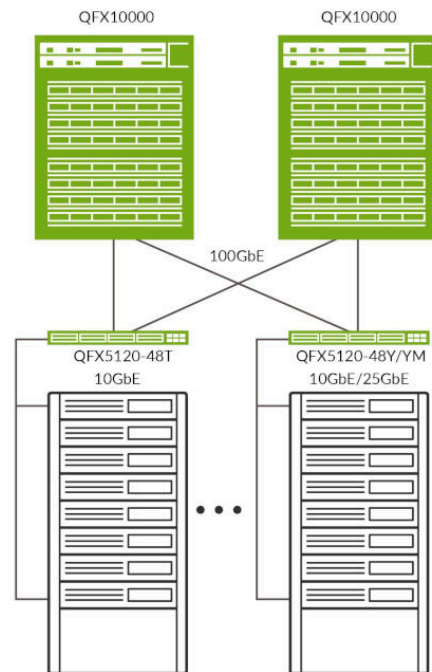


Figure 2: QFX5120-48T and QFX5120-48Y/48YM as leaf and QFX10000 as spine switches

### Data Center Spine

The QFX5120-32C, with 32 ports of 100GbE, can serve as the spine in small to medium-sized enterprise data centers. The QFX5120-32C can be deployed in a three-stage IP Clos with EVPN-VXLAN overlay to support as many as 1500 server access ports. The QFX5120-48Y/QFX5120-48YM and QFX5120-32C can also be used as a collapsed spine data center deployment, offering Ethernet segment identifier-link aggregation group (ESI-LAG) connectivity to the rest of the network infrastructure blocks.

All QFX5120 switches 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 (MC-LAG), the QFX5120 supports active/active server dual-homing and can use full bisectonal bandwidth from server to switch.

Equipped with Junos OS, the QFX5120 supports the most advanced and robust routing capabilities in the industry, 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 L3 underlay for SDN.

## Campus Deployments

Juniper campus fabrics provide a single, standards-based Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) solution that can be deployed in any campus, whether a two-tier network with a collapsed core distribution or a campus-wide system that involves multiple buildings with separate distribution and core layers.

The QFX5120-48Y and QFX5120-48YM are ideal as campus distribution switches with 10GbE/25GbE downlinks and 40GbE/100GbE uplinks supporting technologies like MC-LAG and EVPN multihoming.

The QFX5120-32C is ideal as a campus core switch with 32 ports of 100GbE and support for technologies like campus fabric core-distribution.

The QFX5120-48T supports dual-speed 1GbE/10GbE, also making it a perfect fit for top-of-rack deployments in campus environments for server connectivity.

Juniper campus fabrics support the following validated architectures:

- MC-LAG and EVPN Multihoming (Collapsed Core/Distribution):** A pair of interconnected QFX5120 switches can be deployed to provide EVPN multihoming (ESI-LAG) or multichassis link aggregation (MC-LAG) in a collapsed core/distribution configuration. This eliminates the need for Spanning Tree Protocol (STP) across the campus network by providing multihoming capabilities from the access to the

distribution layer, while distribution to the core is an L3 IP fabric. ESI-LAG also supports horizontal scaling with more than two devices in the distribution layer and can extend EVPN to the core.

- Campus Fabric Core-Distribution:** A pair of interconnected QFX5120 switches can provide EVPN L2 and L3 VXLAN gateway support. This eliminates the need for STP across the campus network by providing a multihoming capability from the access to the distribution layer, while distribution to the core is an L3 IP fabric using EVPN technology. The IP fabric can also extend to connect multiple enterprise buildings, while VXLAN allows stretching of L2 across buildings. An IP Clos network between the distribution and the core layers can exist in two modes, both of which are supported by the QFX5120:
  - Centrally routed bridging overlay:** An IRB interface placed at a central location in the fabric (in this case, a core device)
  - Edge routed bridging overlay:** An IRB interface placed at the edge of the fabric (in this case, a distribution device)
- Campus Fabric IP Clos:** The Campus Fabric IP Clos architecture pushes VXLAN Layer 2/3 gateway functionality to the access layer. In this architecture, the QFX5120 switch acts as an IP fabric distribution switch.

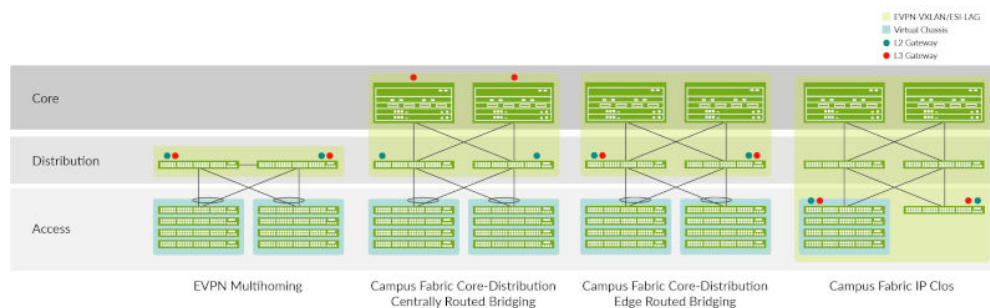


Figure 3: QFX5120 as distribution and core in EVPN multihoming and campus fabric architectures

## Features and Benefits

- **Automation:** The QFX5120 supports a number of network automation and plug-and-play operational features, including ZTP and event scripts, automatic rollback, and Python scripting.
- **Flexible forwarding table:** The QFX5120 includes a unified forwarding table, 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 291,000 MAC addresses. In L3 mode, the QFX5120 table can support 211,000 host entries. 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 Open vSwitch Database (OVSDB) for programming the overlays.
- **MACsec and hop-by-hop encryption:** The QFX5120-48YM supports IEEE 802.1AE MACsec AES-256, providing link-layer data confidentiality, data integrity, and data origin authentication. The MACsec feature enables the QFX5120-48YM to support 2/4 Tbps (uni/bidirectional) of hardware-based traffic encryption on all 100GbE, 40GbE, 25GbE, 10GbE, and 1GbE ports. Defined by IEEE 802.1AE, MACsec provides secure, encrypted communication at the link layer that is capable of identifying and preventing threats from denial-of-service (DoS) and intrusion attacks, as well as man-in-the-middle, masquerading, passive wiretapping, and playback attacks launched from behind the firewall. When MACsec is deployed on switch ports, all traffic is encrypted on the wire, but traffic inside the switch is not. This allows the switch to apply network capabilities such as quality of service (QoS) and

sFlow to each packet without compromising the security of packets on the wire.

In addition, Ethernet-based WAN networks can use MACsec to provide link security over long haul connections. MACsec is transparent to Layer 3 and higher layer protocols and is not limited to IP traffic; it works with any type of wired or wireless traffic carried over Ethernet links.

- **Virtual chassis:** The QFX5120 supports Juniper Networks' unique virtual chassis technology, which enables interconnected switches to operate as a single, logical device with a single IP address. This technology allows campus enterprises to eliminate STP and efficiently utilize network links. QFX5120-48Y (starting with Junos 19.3), QFX5120-32C (starting with Junos 20.3), QFX5120-48T (starting with Junos 20.2), and QFX5120-48YM (starting with Junos 23.1) support virtual chassis feature.

QFX5120-48T, QFX5120-48Y, and QFX5120-32C switch models support HiGig protocol for forming a virtual chassis. However, the QFX5120-48YM switch model supports only HGoE protocol for virtual chassis formation.

## Management, Monitoring, and Analytics

**Data Center Fabric Management:** Juniper Apstra provides operators with the power of intent-based network design to help ensure changes required to enable data center services can be delivered rapidly, accurately, and consistently. Operators can further benefit from the built-in assurance and analytics capabilities to resolve Day 2 operations issues quickly.

Apstra Key Features

- Automated deployment and zero-touch deployment
- Continuous fabric validation
- Fabric lifecycle management
- Troubleshooting using advanced telemetry

For more information on Apstra, see <https://www.juniper.net/us/en/products/network-automation/apstra.html>

## Campus Fabric Management: Juniper Mist Cloud

Juniper Mist Wired Assurance brings cloud management and Mist AI to campus fabrics. It sets a new standard moving away from traditional network management towards AI-driven operations, while delivering better experiences to connected devices.

- Automated deployment and Zero Touch Deployment
- Anomaly detection
- Root cause analysis

For more information on Mist Wired Assurance, see <https://www.juniper.net/us/en/products/cloud-services/wired-assurance.html>

## Paragon Insights (formerly HealthBot)

Combining the power of telemetry, programmability, advanced algorithms, and machine learning, Juniper® Paragon Insights (formerly HealthBot) delivers the following features and benefits for enhanced monitoring and analytics:


- Key performance indicator collection and visualization
- Anomaly detection
- Root cause analysis
- Automated remediation
- Multivendor support
- Customizable playbooks
- JTI telemetry

For more information on Paragon Insights, see <https://www.juniper.net/content/dam/www/assets/datasheets/us/en/network-automation/paragon-insights-datasheet.pdf>.

## Junos Telemetry Interface

The QFX5120 switch 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 paths 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



## QFX5120 Switch Specifications

### Hardware

#### Switching Capacity

- QFX5120-48Y: 2/4 Tbps (uni/bidirectional)/2 Bpps
- QFX5120-48YM: 2/4 Tbps (uni/bidirectional)/2 Bpps
- QFX5120-48T: 1.08/2.16 Tbps (uni/bidirectional)/1 Bpps
- QFX5120-32C: 3.2/6.4 Tbps (uni/bidirectional)/2 Bpps
- Switching mode (all models): Cut-through and store-and-forward

### Weight

- QFX5120-48Y: 23.7 lb (10.75 kg)
- QFX5120-48YM: 23.37 lb (10.60 kg)
- QFX5120-48T: 24.25 lb (11 kg)
- QFX5120-32C: 21.12 lb (9.58 kg)

### Dimensions (H x W x D)

- QFX5120-48Y: 1.72 x 17.36 x 20.48 in. (4.37 x 44.09 x 52.02 cm)
- QFX5120-48YM: 1.72 x 17.36 x 20.48 in. (4.37 x 44.09 x 52.02 cm)
- QFX5120-48T: 1.72 x 17.36 x 20.48 in. (4.37 x 44.09 x 52.02 cm)
- QFX5120-32C: 1.7 x 17.26 x 20.27 in. (4.32 x 43.84 x 51.5 cm)

## Power Consumption

- QFX5120-48Y
  - Max power draw\*: 283 W
  - Typical power draw\*\*: 185 W
- QFX5120-48YM
  - Max power draw\*: 446 W, 460 W with MACsec
  - Typical power draw\*\*: 310 W, 321 W with MACsec
- QFX5120-48T
  - Max power draw\*: 251 W
  - Typical power draw\*\*: 205 W
- QFX5120-32C
  - Max power draw\*\*: 365 W
  - Typical power draw\*\*: 191 W

\* Max power consumption measured at 40°C ambient temperature with SR optics at 100% load with IMIX traffic.  
 \*\* Typical power consumption measured at 25°C ambient temperature with DACs at 50% load with IMIX traffic.  
 Power consumption is subject to operating conditions and unit-to-unit variations.

## Airflow

- Front-to-back (airflow out) for hot aisle deployment
- Back-to-front (airflow in) for cold aisle deployment

## Interface Options

- QFX5120-48Y
  - 2 management ports: 2 x RJ-45 ports
  - 1GbE SFP: 48
  - 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)
  - SFP GbE optical and copper module
  - SFP+ 10GbE optical modules
  - SFP+ direct attach copper (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: Short reach (SR), long reach (LR)
  - QSFP+ to SFP+: 10GbE direct attach breakout copper (1/3 m twinax copper cable)
- QFX5120-48YM
  - 2 management ports: 2 x RJ-45 ports
  - 1GbE SFP: 48
  - 10GbE SFP+: 48/56 (total of 56 ports with breakout cable on port 50 and 52)
  - 25GbE SFP: 48/56 (total of 56 ports with breakout cable on port 50 and 52)

- 40GbE QSFP+: 8 (QSFP+ port 50 and 52 can be configured as a 4 x 10GbE interface or as a 40 Gbps port)
  - 100GbE QSFP28: 8 (QSFP28 port 50 and 52 can be configured as a 4 x 25GbE interface or as a 100 Gbps port)
  - SFP GbE optical
  - 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: Short reach (SR), long reach (LR)
  - QSFP+ to SFP+: 10GbE direct attach breakout copper (1/3 m twinax copper cable)
- QFX5120-48T
    - 1 management port: 1 x RJ-45 port
    - 1GbE RJ45 (copper): 48 (each port is dual speed supporting 1GbE/10GbE)
    - 40GbE QSFP+: 6 (ports 50 and 51 can be configured as a 4 x 10GbE interface or as 40 Gbps interfaces)
    - 100GbE QSFP28: 6 (ports 50 and 51 can be configured as a 4 x 25GbE interface or as 100 Gbps interfaces)
  - QFX5120-32C
    - 1 RJ-45 in-band management port
    - 10GbE SFP+: 2 native ports plus 124 (with 4 x 10GbE breakout cable)
    - 25GbE SFP: 124 (with breakout cable)
    - 40GbE QSFP+: 32 (ports 0-31 can be configured as a 4 x 10GbE interface)
    - 100GbE QSFP28: 32 (ports 0-31 can be configured as a 4 x 25GbE interface)
    - 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+ to SFP+: 10GbE direct attach breakout copper (1/3 m twinax copper cable)

## Common to All Models

- 1 USB 2.0 port
- 1 RS-232 console port
- Supported transceiver and direct attach cable
  - QSFP+ DAC cables: 1/3 m twinax copper
  - QSFP+ optics: SR4, LX4, ESR4, ER4, LR4
  - QSFP28 optics: SR4, ER4, PSM4, CWDM4, LR4

- 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 (QFX5120-48Y/YM and QFX5120-48T) or 5+1 (QFX5120-32C) 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: 4094
- Number of link aggregation groups (LAGs):
  - 80 (QFX5120-48Y/YM, QFX5120-32C)
  - 64 (QFX5120-48T)
- Number of ports per LAG: 64
- IPv4 unicast routes: 2,000,000 prefixes (RIB); 351,000 prefixes (FIB); 211,000 host routes; 64 ECMP paths
- IPv4 multicast routes: 106,000
- IPv6 unicast routes: 500,000 prefixes (RIB); 172,000 prefixes (FIB); 106,000 host routes
- IPv6 multicast routes: 52,000
- Address Resolution Protocol (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
- 100,000 L2 IGMP snooping\*

\*Scale number may vary depending on the test scenarios and other factors.

## 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)
- MACsec with AES256 (QFX5120-48YM only)
- Virtual chassis—up to 2 members on QFX5120-48Y, QFX5120-32C and QFX5120-48T, and up to 4 members on QFX5120-48YM

### Link Aggregation

- MC-LAG
- LAG load sharing algorithm—bridged or routed (unicast or multicast) traffic
- IP: Session Initiation Protocol (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)
- Virtual router
- Dynamic Host Configuration Protocol (DHCP) relay
- Proxy Address Resolution Protocol (ARP)

### EVPN-VXLAN Features

- MAC virtual routing and forwarding (MAC-VRF) multiple EVPN instances (EVI) with service-types vlan-based, vlan-aware, vlan-bundle



- Symmetric inter-IRB routing with anycast gateway and EVPN type-5 instances
- Proxy IGMPv2—EVPN route types 6/7/8
- ARP/ND proxy/suppression
- ESI-LAG A/A multihoming using Enterprise and SP-style interfaces
- Enhanced Ethernet loop detection
- Filter-based forwarding on IRB.VGA
- EVPN advanced route policing
- VLAN-id overlapping using SP-style interfaces
- VLAN rewrite support in EVPN-VXLAN enterprise-style (EP-style) interfaces – vlan-id overlap in EP-style interfaces
- OISM - draft-ietf-bess-evpn-irb-mcast
- VLAN rewrite support in EVPN-VXLAN enterprise-style

### Multicast Features

- Internet Group Management Protocol (IGMP): v1, v2, v3
- IGMP snooping: v1, v2, and v3 (Layer 2 only)
- IGMP filter
- Protocol Independent Multicast-Sparse Mode (PIM-SM), PIM-Source-Specific Multicast (PIM-SSM), PIM-Dense Mode (PIM-DM) in pure IP fabric use case
- 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 in pure IP fabric
- Storm control, port error disable, and autorecovery
- Source MAC address filtering on the port
- DHCP snooping in pure IP fabric use case

### 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 queue (SPQ), shaped-deficit weighted round-robin (SDWRR), weighted random early detection (WRED), weighted tail drop
- 802.1p remarking
- Layer 2 classification criteria: Interface, MAC address, Ethertype, 802.1p, VLAN
- Congestion avoidance capabilities: WRED
- Trust IEEE 802.1p (ingress)
- Remarking of bridged packets
- Default inner to outer DiffServ code point (DSCP) copy for EVPN-VXLAN

### IP Storage

- Priority-based flow control (PFC)—IEEE 802.1Qbb, DCBX
- PFC using DSCP and explicit congestion notification (ECN) for ROCEv2

### 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)

### Management and Analytics Platforms

- Apstra Intent-based System for Data Center
- Juniper Mist Wired Assurance for Campus
- Junos Space® Network Director for Campus
- Paragon Insights

### Device 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
- Inband Flow Analyzer (IFA)
- ZTP
- OpenStack Neutron Plug-in
- Python
- Junos OS event, commit, and OP scripts
- 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 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 1350 Trivial File Transfer Protocol (TFTP) Revision 2
- 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 7348 VXLAN—Virtual extensible Local Area Network
- RFC 8365 NVO—Network Virtualization Overlay Solution Using Ethernet VPN (EVPN-VXLAN)
- OISM - EVPN Optimized Inter-Subnet Multicast (OISM) Forwarding - draft-ietf-bess-evpn-irb-mcast
- IGMP and MLD Proxy for EVPN - draft-ietf-bess-evpn-igmp-ml-d-proxy
- 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 5549 - Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop
- 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
- 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

### 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)

### 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
- IEC 60950-1:2005/A2:2013 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 2015/863/EU



80 Plus Silver PSU Efficiency



Recycled material



Waste Electronics and Electrical Equipment (WEEE)



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



China Restriction of Hazardous Substances (ROHS)

## 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 <https://www.juniper.net/us/en/products.html>.

## Installation and Implementation Service

[Juniper Professional Services](#) offers a Data Center Switching QuickStart program to ensure that the solution is operational and that you have a complete understanding of areas such as configuration and ongoing operations. The QuickStart service provides an onsite consultant who works with your 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-AFI2	QFX5120-48Y, 48x25GbE+8x100GbE, 1 U, AC airflow in
QFX5120-48Y-AFO2	QFX5120-48Y, 48x25GbE+8x100GbE, 1 U, AC airflow out
QFX5120-48Y-DC-AFI2	QFX5120-48Y, 48x25GbE+8x100GbE, 1 U, DC airflow in
QFX5120-48Y-DC-AFO2	QFX5120-48Y, 48x25GbE+8x100GbE, 1 U, DC airflow out
QFX5120-48YM-AFI	48x25GbE+8x100GbE MACsec AES256 switch, AC, back-to-front air flow
QFX5120-48YM-AFO	48x25GbE+8x100GbE MACsec AES256 switch, AC, front-to-back air flow
QFX5120-48YM-DC-AI	48x25GbE+8x100GbE MACsec AES256 switch, DC, back-to-front air flow, DC power
QFX5120-48YM-DC-AO	48x25GbE+8x100GbE MACsec AES256 switch, AC, front-to-back air flow, DC power

Product	Description
QFX5120-48T-AFI	QFX5120-48T, 48x10GbE+6x100GbE, 1 U, AC airflow in
QFX5120-48T-AFO	QFX5120-48T, 48x10GbE+6x100GbE, 1 U, AC airflow out
QFX5120-48T-DC-AFI	QFX5120-48T, 48x25GbE+6x100GbE, 1 U, DC airflow in
QFX5120-48T-DC-AFO	QFX5120-48T, 48x25GbE+6x100GbE, 1 U, DC airflow out
QFX5120-32C-AFI	QFX5120-32C, 32x100GbE, 1 U, AC airflow in
QFX5120-32C-AFO	QFX5120-32C, 32x100GbE, 1 U, AC airflow out
QFX5120-32C-DC-AFI	QFX5120-32C, 32x100GbE, 1 U, DC airflow in
QFX5120-32C-DC-AFO	QFX5120-32C, 32x100GbE, 1 U, DC airflow out
QFX520048Y-APSU-AO	AC Power supply unit, front-to-back airflow for QFX5120-32C-AFO
QFX520048Y-APSU-AI	AC Power supply unit, back-to-front airflow for QFX5120-32C-AFI
QFX520048Y-DPSU-AO	DC Power supply unit, front-to-back airflow for QFX5120-32C-DC-AFO
QFX520048Y-DPSU-AI	DC Power supply unit, back-to-front airflow for QFX5120-32C-DC-AFI
JPSU-650W-AC-AO	AC Power supply unit, front-to-back airflow for QFX5120-48T-AFO and QFX5120-48Y-AFO2
JPSU-650W-AC-AI	AC Power supply unit, back-to-front airflow for QFX5120-48T-AFI and QFX5120-48Y-AFI2
JPSU-650W-DC-AFO	DC Power supply unit, front-to-back airflow for QFX5120-48T-DC-AFO and QFX5120-48Y-DC-AFO2
JPSU-650W-DC-AFI	DC Power supply unit, back-to-front airflow for QFX5120-48T-DC-AFI and QFX5120-48Y-DC-AFI2
JPSU-850W-AC-AFO	AC Power supply unit, front-to-back airflow for QFX5120-48YM-AFO
JPSU-850W-AC-AFI	AC Power supply unit, back-to-front airflow for QFX5120-48YM-AFI
JPSU-850W-DC-AFO	DC Power supply unit, front-to-back airflow for QFX5120-48YM-DC-AO
JPSU-850W-DC-AFI	DC Power supply unit, back-to-front airflow for QFX5120-48YM-DC-AI
QFX520048Y-FAN-AO	front-to-back airflow fan model for QFX5120-32C models
QFX520048Y-FAN-AI	back-to-front airflow fan model for QFX5120-32C models
QFX5110-FANAFO	Front-to-back airflow fan model for QFX5120-48T, QFX5120-48Y, and QFX5120-48YM models
QFX5110-FANAFI	Back-to-front airflow fan model for QFX5120-48T, QFX5120-48Y, and QFX5120-48YM models
QFX512032C-RMK	4 post rack mount kit for QFX5120-32C models
EX-4PST-RMK	4 post rack mount kit for QFX5120-48T, QFX5120-48Y, and QFX5120-48YM models
QFX5K-2PST-RMK	2 post rack mount kit for QFX5120-32C models
EX-RMK	2 post rack mount kit for QFX5120-48T, QFX5120-48Y, and QFX5120-48YM models

**MACsec Encryption Licenses**

S-QFX5K1-MACSEC-1	Class C1 QFX5000, MACsec AES 256 Encryption Sub Software, Term: 1 Yr
S-QFX5K1-MACSEC-3	Class C1 QFX5000, MACsec AES 256 Encryption Sub Software, Term: 3 Yrs
S-QFX5K1-MACSEC-5	Class C1 QFX5000, MACsec AES 256 Encryption Sub Software, Term: 5 Yrs
S-QFX5K1-MACSEC-P	Class C1 QFX5000, MACsec AES 256 Encryption Software, Perpetual

**Flex Software Licenses**

S-QFX5K-C1-A1-3	Flex Sub Software, Class 1 QFX5000 line, Adv 1, Term: 3 Yrs
S-QFX5K-C1-A1-5	Flex Sub Software, Class 1 QFX5000 line, Adv 1, Term: 5 Yrs
S-QFX5K-C1-A1-P	Flex Software, Class 1 QFX5000 line, Adv 1, Perpetual
S-QFX5K-C1-A2-3	Flex Sub Software, Class 1 QFX5000 line, Adv 2, Term: 3 Yrs
S-QFX5K-C1-A2-5	Flex Sub Software, Class 1 QFX5000 line, Adv 2, Term: 5 Yrs

Product	Description
S-QFX5K-C1-A2-P	Flex Software, Class 1 QFX5000 line, Adv 2, Perpetual
S-QFX5K-C1-P1-3	Flex Sub Software, Class 1 QFX5000 line, Prem 1, Term: 3 Yrs
S-QFX5K-C1-P1-5	Flex Sub Software, Class 1 QFX5000 line, Prem 1, Term: 5 Yrs
S-QFX5K-C1-P1-P	Flex Software, Class 1 QFX5000 line, Prem 1, Perpetual
S-QFX5K-C2-A1-3	Flex Sub Software, Class 2 QFX5000 line, Adv 1, Term: 3 Yrs
S-QFX5K-C2-A1-5	Flex Sub Software, Class 2 QFX5000 line, Adv 1, Term: 5 Yrs
S-QFX5K-C2-A1-P	Flex Software, Class 2 QFX5000 line, Adv 1, Perpetual
S-QFX5K-C2-A2-3	Flex Sub Software, Class 2 QFX5000 line, Adv 2, Term: 3 Yrs
S-QFX5K-C2-A2-5	Flex Sub Software, Class 2 QFX5000 line, Adv 2, Term: 5 Yrs
S-QFX5K-C2-A2-P	Flex Software, Class 2 QFX5000 line, Adv 2, Perpetual
S-QFX5K-C2-P1-3	Flex Sub Software, Class 2 QFX5000 line, Prem 1, Term: 3 Yrs
S-QFX5K-C2-P1-5	Flex Sub Software, Class 2 QFX5000 line, Prem 1, Term: 5 Yrs
S-QFX5K-C2-P1-P	Flex Software, Class 2 QFX5000 line, Prem 1, Perpetual

**Services Licenses**

SVC-COR-QFX51-48M	Juniper Care Core Support for QFX5120-48YM Switches
SVC-ND-QFX51-48M	Juniper Care Next Day Support for QFX5120-48YM
SVC-COR-QFX5KC1-MS	Juniper Care Core Support for S-QFX5KC1-MACSEC-P

**Advanced Telemetry Feature (ATF) Licenses**

S-QFX5K-C1-ATF-P	SW, QFX5K, Class 1, ATF License SKU, Perpetual, SVC Customer Support to be purchased separately
S-QFX5K-C2-ATF-P	SW, QFX5K, Class 2, ATF License SKU, Perpetual, SVC Customer Support to be purchased separately

Note: an ATF license is needed on each Inband Flow Analyzer (IFA) terminating node. For definition of IFA terminating node and models supported, please refer to our [Monitoring, Sampling, and Services Interfaces User Guide](#).

## Optics and Transceivers

Up-to-date information on supported optics for the QFX5120 switches can be found in the Hardware Compatibility Tool : [QFX5120-32C](#) , [QFX5120-48T](#) , [QFX5120-48YM](#) , and [QFX5120-48Y](#).

## About Juniper Networks

Juniper Networks believes that connectivity is not the same as experiencing a great connection. Juniper's AI-Native Networking Platform is built from the ground up to leverage AI to deliver the best and most secure user experiences from the edge to the data center and cloud. Additional information can be found at Juniper Networks ([www.juniper.net](http://www.juniper.net)) or connect with Juniper on X (Twitter), LinkedIn, and Facebook.

### 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**

**[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: +31.207.125.700**

