


# QFX5700 LINE OF SWITCHES DATASHEET



## Product overview

The [QFX5700 and QFX5700E Switches](#), which are high-density, cost-optimized, 5 U 400GbE, 8 slot fabric-less modular platforms, are ideal for [data centers](#) and campus distribution/core networks where capacity and cloud services are being added as business needs grow. These services require higher network bandwidth per rack, as well as flexibility.

The line supports:

- spine-and-leaf deployments in [enterprise](#), [service provider](#), and [cloud provider data centers](#)
- evolving business and network needs
- deployment versatility

## Product description

The Juniper Networks® QFX5700 and QFX5700E Switches are next-generation, modular and fabric-less spine-and-leaf switches that offer flexibility, cost efficiency with lower-per-bit, high-density [400GbE](#), 100GbE, 50GbE, 40GbE, 25GbE, and 10GbE interfaces for server and intra-fabric connectivity.

A versatile, future-proven solution for today's data centers, the QFX5700 and QFX5700E switches support and deliver a diverse set of use cases. The line supports advanced Layer 2, Layer 3, and Ethernet VPN (EVPN)-[Virtual Extensible LAN \(VXLAN\)](#) features. For large public cloud providers—early adopters of high-performance servers to meet explosive workload growth—the QFX5700 and QFX5700E switches support very large, dense, and fast 400GbE IP fabrics based on proven Internet scale technology. For enterprise data center customers seeking investment protection as they transition their server farms from 10GbE to 25GbE, the switches also provide a high radix-native 100GbE/400GbE EVPN-VXLAN spine option at reduced power and a smaller footprint.

The QFX5700 and QFX5700E switches support diverse use cases such as Data Center Fabric Spine, EVPN-VXLAN Fabric, Data Center Interconnect/Border, Secure DCI, multi-tier campus, campus fabric, and connecting firewall clusters in the DC to the fabric. Delivering 25.6 Tbps of bidirectional bandwidth, the switches are optimally designed for spine-and-leaf deployments in enterprise, high-performance computing (HPC), service provider, and cloud data centers.

The QFX5700 and QFX5700E switches are a modular merchant silicon-based chassis offering a variety of port configurations, including 400GbE, 100GbE, 50GbE, 40GbE, 25GbE, and 10GbE. The switches are equipped with up to four AC or DC power supplies, providing N+N feed redundancy or N+1 PSU redundancy when all power supplies are present. Two hot-swappable fan trays offer front-to-back (AFO) airflow, providing N+1 fan rotor redundancy at chassis.

The QFX5700 and QFX5700E switches include an Intel Hewitt Lake 6 core to drive the control plane, which runs the [Junos® OS Evolved](#) operating system software.

## Product highlights

The QFX5700 and QFX5700E switches include the following capabilities. Please refer to the Specifications section for current shipping features.

### Native 400GbE configuration

The QFX5700 & QFX5700E switches offer 32 ports of 400GbE in a modular 8 slot 5 U form factor.

For a complete list of features, refer to the Specifications section and check out Feature Explorer.

### High-density configurations

The QFX5700 and QFX5700E switches are optimized for high-density fabric deployments, providing options for 32 ports of 400GbE, 64 ports of 200GbE (using Breakout cable), 128 ports of 100GbE, and 40GbE, 144 ports of 50GbE/ 40GbE/ 25GbE/ 10GbE with the opportunity to scale-as-you-grow.

### Flexible connectivity options

The QFX5700 and QFX5700E switches offer a choice of interface speeds for server and intra-fabric connectivity, providing deployment versatility and investment protection.

## Key product differentiators

### Increased scale and buffer

The QFX5700 switch provides enhanced scale with up to 1.24 million routes, 80,000 firewall filters, and 160,000 media access control (MAC) addresses. It supports high numbers of egress IPv4/IPv6 rules by programming matches in egress ternary content addressable memory (TCAM) along with ingress TCAM.

### 132MB shared packet buffer

Today's cloud-native applications have critical dependency on buffer size to prevent congestion and packet drops. The QFX5700 and QFX5700E switches have 132 MB shared packet buffer that is allocated dynamically to congested ports.

## Features and benefits

- Automation and programmability:** The QFX5700 and QFX5700E switches support several network automation features for plug-and-play operations, including zero-touch provisioning (ZTP), Network Configuration Protocol (NETCONF), Juniper Extension Toolkit (JET), Junos telemetry interface, operations and event scripts, automation rollback, and Python scripting.
- Cloud-level scale and performance:** QFX5700 and QFX5700E switches support best-in-class cloud-scale L2/L3 deployments with a latency as low as 900 ns and superior scale and performance. The switches support up to 128 link aggregation groups, 4096 VLANs, and Jumbo frames of 9216 bytes. Junos OS Evolved provides configurable options through a CLI, enabling each QFX5700 and QFX5700E to be optimized for different deployment scenarios.
- VXLAN overlays:** The QFX5700 and QFX5700E switches are 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, which can operate without a controller or be orchestrated with an SDN controller.
- RoCEv2:** As a switch capable of transporting data as well as storage traffic over Ethernet, the QFX5700 and QFX5700E switches provide an IEEE data center bridging (DCB) converged network between servers with disaggregated flash storage arrays or an NVMe-enabled storage-area network (SAN). The switches offer a full-featured DCB implementation that provides strong monitoring capabilities on the top-of-rack switch for SAN and LAN administration teams to maintain clear separation of management. The RDMA over Converged Ethernet version 2 (RoCEv2) transit switch functionality, including priority-based flow control (PFC) and Data Center Bridging Capability Exchange (DCBX), are included as part of the default software.
- Junos Evolved features:** The QFX5700 and QFX5700E switches support features such as L2/L3 unicast, EVPN-VXLAN, BGP add-path, RoCEv2 and congestion management, multicast, 128-way ECMP, dynamic load balancing capabilities, enhanced firewall capabilities, and monitoring.
- Junos OS Evolved architecture:** Junos OS Evolved is a native Linux operating system that incorporates a modular design of independent functional components and enables individual components to be upgraded independently while the system remains operational. Component failures are localized to the specific component involved and can be corrected by upgrading and restarting that specific component without having to bring down the entire device. The switches control and data plane processes can run in parallel, maximizing CPU utilization, providing support for containerization, and enabling application deployment using LXC or Docker.
- Retained state:** State is the retained information or status pertaining to physical and logical entities. It includes both operational and configuration state, comprising committed

configuration, interface state, routes, hardware state, and what is held in a central database called the distributed data store (DDS). State information remains persistent, is shared across the system, and is supplied during restarts.

- **Feature support:** All key networking functions such as routing, bridging, management software, and management plane interfaces, as well as APIs such as CLI, NETCONF, JET, Junos telemetry interface, and the underlying data models, resemble those supported by the Junos operating system. This ensures compatibility and eases the transition to Junos Evolved.

## Deployment options

### Data center fabric deployments

The QFX5700 and QFX5700E switches can be deployed as a universal device in cloud data centers to support 100GbE server access and 400GbE spine-and leaf configurations, optimizing data center operations by using a single device across multiple network layers (see Figure 1). The switches can also be deployed in more advanced overlay architectures like an EVPN-VXLAN fabric. Depending on where tunnel terminations are desired, the switches can be deployed in the EVPN-VXLAN Edge Routed Bridging architecture or in Bridged Overlay (BO).

Juniper offers complete flexibility and a range of data center fabric designs that cater to data centers of different sizes, scale built by Cloud Operators, Service Providers and Enterprises. Here're the main data center fabric design options where QFX5700 and QFX5700E switches can be used for high port density server leaf, spine node or border-leaf node:

- Architecture 1: **ERB - Edge Routed Bridging EVPN-VXLAN** with distributed anycast IP gateway architecture supporting L2 and L3 for Enterprises and 5G Telco-Cloud. This type of design offers a combination of L2 stretch between multiple leaf/ToR switches and L2 active/active multihoming to the server with MAC-VRF EVI L2 virtualization support as well as L3 IP VRF virtualization at the leaf/ToR through the Type-5 EVPN-VXLAN. This type of design in DC use-case can be used to connect in redundant and optimized way the servers/compute nodes, Blade Centers, IP storage nodes running ROCEv2, as well as other appliances.
- Architecture 2: **BO - Bridged Overlay EVPN-VXLAN design** using MAC-VRF instances and different EVPN service-types (vlan-aware, vlan-bundle, vlan-based). In this case an external to the fabric first hop IP gateway can be used – for example at the firewall or external existing DC gateway routers. In this design the DC fabric is offering L2 active/active multihoming using ESI-LAG and fabric wide L2 stretch between the leaf ToR nodes.
- Architecture 3: **Seamless Data Center Interconnect (DCI) for ERB fabric design** – DCI border-leaf design with seamless T2/T2 EVPN-VXLAN to EVPN-VXLAN tunnel stitching (RFC 9014) and T5/T5 EVPN-VXLAN tunnel stitching support. With this design, the data center benefits from geographical redundancy for the application deployed in private cloud DC. In this case, the QFX5700 and QFX5700E switches can also be used as a border-leaf node.
- Architecture 4: **Collapsed Spine design with ESI-LAG support and anycast IP** – in this case the pair of QFX5700 and QFX5700E switches is deployed with a back-to-back connect, without spine layer. The L2 active/active multihoming using ESI-LAG is used for the server NIC high availability as well as anycast IP gateway.

### 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 are:

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

For more information on Apstra, see [www.juniper.net/us/en/products/network-automation/apstra/apstra-system.html](http://www.juniper.net/us/en/products/network-automation/apstra/apstra-system.html).

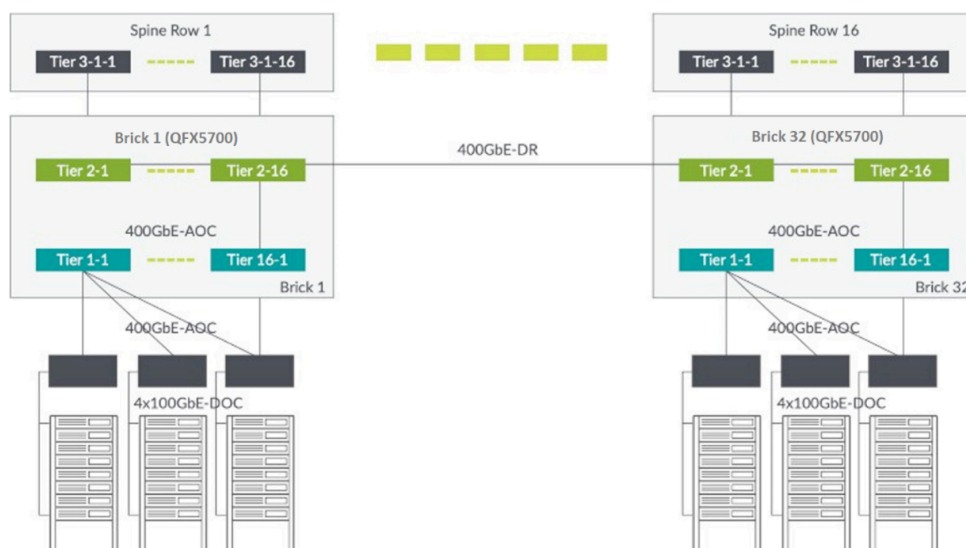


Figure 1: Typical cloud data center deployment for the QFX5700 and QFX5700E

## Campus fabric deployments

### EVPN-VXLAN for Campus Core, Distribution, and Access

The QFX5700 and QFX5700E switches can be deployed in campus distribution/core layer networks using 10GbE/25GbE/40GbE/100GbE ports to support technologies such as EVPN multihoming and campus Fabric.

Juniper offers complete flexibility in choosing any of the following validated EVPN-VXLAN designs that cater to networks of different sizes, scale, and segmentation requirements:

- EVPN multihoming (collapsed core or distribution):** A collapsed core architecture combines the core and distribution layers into a single switch, turning the traditional three-tier hierarchical network into a two-tier network. EVPN Multihoming on a collapsed core eliminates the need for Spanning Tree Protocol (STP) across campus networks by providing link aggregation capabilities from the access layer to the core layer. This topology is best suited for small to medium distributed enterprise networks and allows for consistent VLANs across the network. This topology uses ESI (Ethernet Segment Identifier) LAG (Link Aggregation) and is a standards-based protocol.
- Campus fabric core distribution:** When EVPN VXLAN is configured across core and distribution layers, it becomes a campus Fabric Core Distribution architecture, which can be configured in two modes: centrally or edge routed bridging overlay. This architecture provides an opportunity for an

administrator to move towards campus-fabric IP Clos without fork-lift upgrade of all access switches in the existing network, while bringing in the advantages of moving to a campus fabric and providing an easy way to scale out the network.

- Campus fabric IP Clos:** When EVPN VXLAN is configured on all layers including access, it is called the campus fabric IP Clos architecture. This model is also referred to as “end-to-end,” given that VXLAN tunnels are terminated at the access layer.

In all these EVPN-VXLAN deployment modes, QFX5700 and QFX5700E switches can be used in Distribution or core as seen in Figure 2. All three topologies are standards-based and hence are inter-operable with 3rd party vendors.

### Managing AI-Native campus fabric with the Juniper Mist Cloud

Juniper Mist Wired Assurance brings cloud management and Mist AI to campus fabric. It sets a new standard that moves away from traditional network management towards AI-driven operations, while delivering better experiences to connected devices. The Juniper Mist Cloud streamlines deployment and management of campus fabric architectures by allowing:

- Automated deployment and zero touch deployment (ZTD)
- Anomaly detection
- Root cause analysis

For more information, read the [Juniper Mist Wired Assurance datasheet](#).




Figure 2: Campus fabrics showing Virtual Chassis and EVPN-VXLAN architectures

### Architecture and key components

The QFX5700 and QFX5700E switches can be used in L2 fabrics and L3 networks. You can choose the architecture that best suits your deployment needs and easily adapt and evolve as requirements change over time. The QFX5700 and QFX5700E switches serve as the universal building block for these switching architectures, enabling data center operators to build cloud networks in their own way.

**Layer 3 fabric:** For customers looking to build scale-out data centers, a Layer 3 spine-and-leaf Clos fabric provides predictable, nonblocking performance and scale characteristics. A two-tier fabric built with QFX5700 and QFX5700E switches as leaf devices and [Juniper Networks QFX10000 modular switches](#) in the spine can scale to support up to 128 40GbE ports or 128 25GbE and/or 10GbE server ports in a single fabric.

Junos OS Evolved ensures a high feature and bug fix velocity and provides first-class access to system state, allowing customers to run DevOps tools, containerized applications, management agents, specialized telemetry agents, and more.

### Junos telemetry interface

The QFX5700 and QFX5700E switches support Junos telemetry interface, a modern telemetry streaming tool that provides performance monitoring in complex, dynamic data centers. Streaming data to a performance management system lets network administrators measure trends in link and node utilization and troubleshoot issues such as network congestion in real time.


Junos telemetry interface provides:

- Application visibility and performance management by provisioning sensors to collect and stream data and analyze the 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 correlating overlay and underlay networks

### Power consumption

Parameters	8x QFX5K-FPC-20Y 144x50G	8x QFX5K-FPC-16C 128x100G	8x QFX5K-FPC-4CD 32x400G
Typical power draw*	751 W	1,259 W	1,095 W
Maximum power draw**	1,622 W	2,271 W	1,762 W

\*Typical power consumption is measured at 25°C ambient temperature, and at 50% load with IMIX traffic, without MACsec, excludes transceivers.  
 \*\*Max power consumption is measured at 40°C ambient temperature with SR optics, and at 100% load with IMIX traffic, with MACsec on QFX5K-FPC-20Y and QFX5K-FPC-16C, without MACsec on QFX5K-FPC-4CD.  
 Power consumption is subject to operating conditions and unit-to-unit variations; the power consumption measurements were taken at 200-277 VAC PSUs.



## Specifications

### Hardware

Table 1. QFX5700 and QFX5700E System Capacity

Specification	QFX5700
<b>System throughput</b>	Up to 12.8 / 25.6 Tbps (uni / bidirectional)
<b>Forwarding capacity</b>	5.3 Bpps
<b>Port density without breakout</b>	32 ports of QSFP56-DD (400GbE); 128 ports of QSFP28 (100GbE) or QSFP+ (40GbE); or 144 ports of SFP56 (50GbE), SFP28 (25GbE), or SFP+ (10GbE)
<b>Max ports with breakouts</b>	Line card QFX5K-FPC-16C: 32 4x25GbE, 32 4x10GbE Line card QFX5K-FPC-4CD: 32 4x100GbE, 32 4x25GbE, 32 4x10GbE Go to the Port Checker Tool to see different port combinations for line cards: QFX5K-FPC-20Y, QFX5K-FPC-16C, & QFX5K-FPC-4CD

Table 2. QFX5700 and QFX5700E system specifications

Specification	QFX5700
<b>Dimensions (W x H x D)</b>	19.0 in x 8.74 in (5RU) x 32 in (48.2 x 22.2 x 81.5 cm)
<b>Rack units</b>	5 U
<b>Weight</b>	153.8 lbs. (69.8 kg) with all FRUs installed
<b>Operating system</b>	Junos OS Evolved
<b>CPU</b>	Intel Hewitt Lake, 32GB DDRAM
<b>Power</b>	<ul style="list-style-type: none"> <li>Redundant (N+N) hot-pluggable 3000 W AC/ DC power supplies (2n)</li> </ul>
<b>Cooling</b>	<ul style="list-style-type: none"> <li>Two hot-pluggable Fan trays with Four Counter Rotating Fans in each Fan Tray</li> <li>N+1 Fan rotor redundancy at Chassis level</li> <li>Front to Back air cooling</li> </ul>
<b>Total packet buffer</b>	132MB
<b>Warranty</b>	Juniper standard one-year warranty

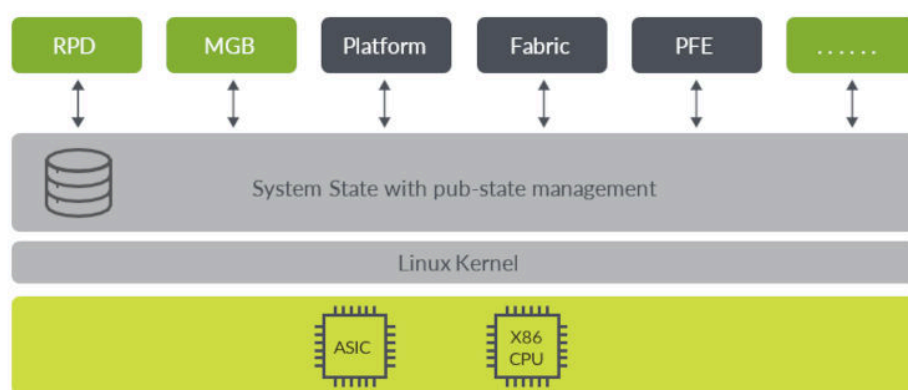


Figure 3: Cloud/Carrier-Class Junos OS Evolved Network Operating System

### Software scale

Software	QFX5700	QFX5700E
Operating System	Junos Evolved	Junos Evolved
MAC addresses per system	160,000	96,000
VLAN IDs	4,000	4,000
Number of link aggregation groups (LAGs)	128	128
Ingress routed ACL (RACL)	20K	10,000 terms configured to one filter
Ingress VLAN ACL (VACL)	20K	10,000 terms configured to one filter
Ingress port ACL (PACL)	20K	10,000 terms configured to one filter
Egress routed ACL (RACL)	1,000	1,000 terms configured to one filter
Egress VLAN ACL (VACL)	2,000	2,000 terms configured to one filter
Egress port ACL (PACL)	2,000	2,000 terms configured to one filter
unicast routes IPv4/v6	1.24 million / 610,000 routes	732,000 / 428,000 routes
Host routes ipv4/ipv6	160,000 / 80,000	96,000 / 48,000
ARP entrees	32,000 (tunnel mode) 64,000 (non-tunnel mode)	32,000 (tunnel mode) 64,000 (non-tunnel mode)

Software	QFX5700	QFX5700E
Jumbo frame	9,216 Bytes	9,216 Bytes
Traffic mirroring destination ports per switch	64	64
Maximum number of mirroring sessions	4	4
Traffic mirroring destination vlans per switch	60	60
Shared Packet Buffer (MB)	132	132
Firewall filters	80,000	80,000

- Number of ports per LAG: 64
- Neighbor Discovery Protocol (NDP) entries: 32,000 (tunnel mode); 64,000 (non-tunnel mode)
- Generic routing encapsulation (GRE) tunnels: 1000
- Traffic mirroring: 8 destination ports per switch

### 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 VLAN Spanning Tree Protocol (VSTP) running concurrently
- VLAN—IEEE 802.1Q VLAN trunking
- Routed VLAN interface (RVI)
- Port-based VLAN
- Static MAC address assignment for interface
- MAC learning disable
- Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)

### Link aggregation

- 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
  - L2 and non-IP: MAC SA, MAC DA, Ether type, VLAN ID, source port

### Layer 3 features

- Static routing
- OSPF v1/v2
- OSPF v3
- Filter-based forwarding
- Virtual Router Redundancy Protocol (VRRP)
- IPv6
- Virtual routers
- Loop-free alternate (LFA)
- BGP (Advanced Services or Premium Services license)
- IS-IS (Advanced Services or Premium Services license)
- Dynamic Host Configuration Protocol (DHCP) v4/v6 relay
- VR-aware DHCP
- IPv4/IPv6 over GRE tunnels (interface-based with decap/encap only)

### Multicast

- Internet Group Management Protocol (IGMP) v1/v2
- Multicast Listener Discovery (MLD) v1/v2
- IGMP proxy, querier
- IGMP v1/v2/v3 snooping

- MLD snooping
- Protocol Independent Multicast PIM-SM, PIM-SSM, PIM-DM

### Security and filters

- Secure interface login and password
- Secure boot
- RADIUS
- TACACS+
- Ingress and egress filters: Allow and deny, port filters, VLAN filters, and routed filters, including management port filters and loopback filters for control plane protection
- Filter actions: Logging, system logging, reject, mirror to an interface, counters, assign forwarding class, permit, drop, police, mark
- SSH v1, v2
- Static ARP support
- Storm control, port error disable, and autorecovery
- Control plane denial-of-service (DoS) protection
- Image rollback

### Quality of Service (QoS)

- L2 and L3 QoS: Classification, rewrite, queuing
- Rate limiting:
  - Ingress policing: 1 rate 2 color, 2 rate 3 color
  - Egress policing: Policer, policer mark down action
  - Egress shaping: Per queue, per port
- 12 hardware queues per port (8 unicast and 4 multicast)
- Strict priority queuing (LLQ), shaped-deficit weighted round-robin (SDWRR), weighted random early detection (WRED)
- 802.1p remarking
- Layer 2 classification criteria: Interface, MAC address, Ether type, 802.1p, VLAN
- Congestion avoidance capabilities: WRED
- Trust IEEE 802.1p (ingress)
- Remarking of bridged packets

### EVPN-VXLAN

- EVPN support with VXLAN transport
- All-active multihoming support for EVPN-VXLAN (ESI-LAG aka EVPN-LAG)
- MAC-VRF (EVI) multiple EVPN service-type support: vlan-based, vlan-aware, vlan-bundle
- ARP/ND suppression aka proxy-arp/nd
- Ingress multicast Replication
- IGMPv2 snooping support fabric wide: using EVPN route type-6

- IGMPv2 snooping support for L2 multihoming scenarios: EVPN route type-7 and type-8
- IP prefix advertisement using EVPN with VxLAN encapsulation

### Data center bridging technologies

- Explicit congestion notification (ECN)
- Data Center Bridging Quantized Congestion Notification (DCQCN)
- Priority-based flow control (PFC)—IEEE 802.1Qbb
- Priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic
- Remote Direct Memory Access (RDMA) over converged Ethernet version 2 (RoCEv2)

### High availability

- Bidirectional Forwarding Detection (BFD)

### Visibility and analytics

- Switched Port Analyzer (SPAN)
- Remote SPAN (RSPAN)
- Encapsulated Remote SPAN (ERSPAN)
- sFlow v5
- Junos telemetry interface

### Management and operations

- Role-based CLI management and access
- CLI via console, telnet, or SSH
- Extended ping and traceroute
- Junos OS Evolved configuration rescue and rollback
- SNMP v1/v2/v3
- Junos OS Evolved XML management protocol
- High frequency statistics collection
- Automation and orchestration
- Zero-touch provisioning (ZTP)
- Python
- Junos OS Evolved event, commit, and OP scripts
- Juniper Apstra Management, Monitoring, and Analytics for Data Center Fabrics
- Juniper Mist Wired Assurance for Campus

### Standards compliance

#### IEEE Standards

- IEEE 802.1D
- IEEE 802.1w

- IEEE 802.1
- IEEE 802.1Q
- IEEE 802.1p
- IEEE 802.1ad
- IEEE 802.3ad
- IEEE 802.1AB
- IEEE 802.3x
- IEEE 802.1Qbb
- IEEE 802.1Qaz
- T11 Standards
- INCITS T11 FC-BB-5

### Environmental ranges

Parameters	QFX5700
Operating temperature	32° F to 104° F (0° C to 40° C)
Storage temperature	-40° F through 158° F (-40° C to 70° C)
Operating altitude	Up to 6000 feet (1828.8 meters)
Relative humidity operating	5 to 90% (noncondensing)
Relative humidity nonoperating	5 to 95% (noncondensing)
Seismic	Designed to meet GR-63, Zone 4 earthquake requirements

### Customer-specific requirements

- GR-1089-Core, Issue 8
- Juniper Inductive GND (IGS)
- Deutsche Telekom (DT) 1TR9
- British Telecommunications (BT) GS7

### 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
- UL 62368-1 Second Edition
- UL IEC 62328-1 Second Edition

### Electromagnetic compatibility global

- FCC 47 CFR Part 15
- ICES-003 / ICES-GEN
- BS EN 55032



- BS EN 55035
- EN 300 386 V1.6.1
- EN 300 386 V2.1.1
- BS EN 300 386
- EN 55032
- CISPR 32
- EN 55035
- CISPR 35
- IEC/EN 61000 Series
- IEC/EN 61000-3-2
- IEC/EN 61000-3-3
- AS/NZS CISPR 32
- VCCI-CISPR 32
- BSMI CNS 15936
- KS C 9835 (Old KN 35)
- KS C 9832 (Old KN 32)
- KS C 9610
- BS EN 61000 Series

## Telco

- Common Language Equipment Identifier (CLEI) code

## Environmental compliance



Restriction of Hazardous Substances (ROHS) 6/6



China Restriction of Hazardous Substances (ROHS)



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



Waste Electronics and Electrical Equipment (WEEE)



Recycled material



80 Plus Silver PSU Efficiency

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

## Ordering information

Product Number	Description
<b>QFX5700/QFX5700E hardware</b>	
QFX5700-CHAS	QFX5700/QFX5700E spare chassis
QFX5700-BASE-AC	QFX5700 (hardware only; software services sold separately), with 1 5700-FEB 1 RCB, redundant fans, 2 AC power supplies, front- to-back airflow
QFX5700-BASE-DC	QFX5700 (hardware only, software services sold separately), with 1 5700-FEB, 1 RCB, redundant fans, 2 DC power supplies, front- to-back airflow
QFX5700E-BASE-AC	QFX5700E (hardware only; software services sold separately), with 1 5700E-FEB, 1 RCB, redundant fans, 2 AC power supplies, front- to-back airflow
QFX5700E-BASE-DC	QFX5700E (hardware only; software services sold separately), with 1 5700E-FEB, 1 RCB, redundant fans, 2 DC power supplies, front- to-back airflow
<b>QFX5700/QFX5700E line cards</b>	
QFX5K-FPC-4CD	4X400G line card for QFX5700/QFX5700E chassis
QFX5K-FPC-20Y	10G/25G(SFP) line card for QFX5700/QFX5700E chassis
QFX5K-FPC-16C	16X100G line card for QFX5700/QFX5700E chassis
<b>QFX5700/QFX5700E Power Supply</b>	
JNP-3000W-AC-AFO	AC PS 3000W, AFO
JNP-3000W-DC-AFO	DC PS 3000W, AFO
<b>Software Licenses SKUs</b>	
S-QFX5KC3-MACSEC-3	MACsec Software feature license for QFX5700/QFX5700E, 16(100G) ports + 4(400G) ports + 20(10G/25G) ports, 3 Year
S-QFX5KC3-MACSEC-5	MACsec Software feature license for QFX5700/QFX5700E, 16(100G) ports + 4(400G) ports + 20(10G/25G) ports, 5 Year
S-QFX5KC3-MACSEC-P	MACsec Software feature license for QFX5700/QFX5700E, 16(100G) ports + 4(400G) ports + 20(10G/25G) ports, Perpetual
S-QFX5K-C3-A1-X (X=3,5)	Base L3 Software Subscription (X Years; X=3,5) License for QFX5700/QFX5700E
S-QFX5K-C3-A2-X (X=3,5)	Advanced Software Subscription (X Years; X=3,5) License for QFX5700/QFX5700E
S-QFX5K-C3-P1-X (X=3,5)	Premium Software Subscription (X Years; X=3,5) License for QFX5700/QFX5700E
<b>Cable SKUs</b>	
CBL-JNP-SDG4-JPL	Cable Specific, Japan
CBL-JNP-SDG4-TW	Cable Specific, Taiwan
CBL-JNP-SDG4-US-L6	Cable Specific, US/North America, L6
CBL-JNP-PWR-EU	Cable Specific, EU,Africa,China
CBL-JNP-SDG4-US-L7	Cable Specific, US/North America, L7
CBL-JNP-SDG4-IN	Cable Specific, India
CBL-JNP-SDG4-SK	Cable Specific, South Korea
<b>Additional SKUs</b>	
JNP5K-FEB-BLNK	Blank cover for empty FEB slot
JNP5K-FPC-BLNK	Blank cover for empty FPC (Line card) slot
JNP5K-RCB-BLNK	Blank cover for empty RCB (Routing Control Board) slot
JNP5K-RMK-4POST	4-Post Rack Mount Kit
QFX5K-EMI	Cable Manager
JNP5700-FAN	Airflow out (AFO) front-to-back airflow fans for QFX5700/QFX5700E

## Optics and transceivers

The QFX5700 and QFX5700E switches support varying port speeds at 400G, 100G, 50G, 40G, 25G, and 10G with different transceiver options of direct attach copper cables, active optical cables, break out cables (DACBO and AOCBO). Up-to-date information on supported optics can be found on the Hardware Compatibility Tool at <https://apps.juniper.net/hct/product/>.

## 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**

