



* Product configuration and appearance may vary

JUNIPER QFX5240 LINE OF SWITCHES



Product description

Continuous evolution of AI/ML technology along with new applications are driving the next major shift in bandwidth requirement within the [data center fabric](#). Juniper QFX5240 line of Switches is a next-generation, fixed-configuration platform designed for spine, leaf, and border switch roles. The switch provides flexible, cost-effective, high-density 800GbE, 400GbE, 100GbE, and 50GbE interfaces for intra-IP fabric connectivity as well as higher density 200/400GbE NIC connectivity for AI/ML use cases. It's 51.2 Tbps unidirectional throughput meets the bandwidth requirement of AI/ML workloads and storage systems with latency starting at 750 ns (Cut-through) and 950 ns (Store and Forward). Remote direct memory access (RDMA) is the de-facto data transfer technology used in AI/ML workloads, and it uses remote direct memory access over Converged Ethernet v2 (ROCEv2) for transport at the network layer. The QFX5240 line of switches supports ROCEv2 along with congestion management features like Priority Flow Control (PFC), explicit congestion notification (ECN), and data center quantized congestion notification (DCQCN).

The QFX5240 line of switches helps reduce the number of network nodes deployed—decreasing the total power consumption of the data center fabric and improving the carbon footprint of the data center. These improvements are possible by having different breakout options like 128x400GbE, 256x200GbE, and 320x100GbE.

Product overview

The [QFX5240 line of switches](#) meet the advanced [AI data center networking](#) requirements of large-scale clusters. The QFX5240 line works with the automation—such as [Apstra Data Center Director](#) (formerly Juniper Apstra)—to assure daily operation in AI and ML workload training and access.

The QFX5240 line of switches:

- Deliver high-density [800GbE](#) ports on a fixed form factor with software to provide advanced network services tuned to the specific needs of AI/ML workloads
- Are a foundation of AI networks, and their low latencies ensure fast job completion time (JCT) to speed training through high GPU utilization
- Help teams managing AI/ML environments realize improved economics

Table 1. QFX5240 line of switches product highlights

Line of switches		
QFX5241-32OD		QFX5240-64OD QFX5240-64QD QFX5241-64OD QFX5241-64QD
AI data center	– Leaf/spine in AI/ML cluster – ROCEv2 for AI/ML workloads – DCQCN-PFC, ECN for congestion management – Support for PFC watchdog for storm avoidance – Dynamic load balancing (DLB) for better load balancing – Configurable hash-bucket size to suit different flow scale – Selective-DLB—enable/disable DLB based on opcode/byte match	
Cloud-ready data center	– Leaf/spine in IP fabric – Leaf/spine/super spine in EVPN-VXLAN fabric – Support for EVPN-VXLAN – 136K MAC scale – IPv4 unicast routes: • 2,000,000 routes (RIB), 880,000 LPM routes (FIB) and 880,000 host routes (FIB) - IPv6 unicast routes: 2,000,000 routes (RIB), 875,000 LPM routes (FIB) and 875,000 host routes (FIB)	
Port options	– 32 x 800GbE ports – 2 x SFP+ 10GbE ports – 64 x 400GbE (with breakout cable) – 128 x 200GbE (with breakout cable) – 160 x 100GbE (with breakout cable)	– 64 x 800GbE ports – 2 x SFP+ 10GbE ports – 128 x 400GbE (with breakout cable) – 256 x 200GbE (with breakout cable) – 320 x 100GbE (with breakout cable)
Platform parameters	– Throughput: 25.6 Tbps unidirectional – Buffer: 83 MB – Tool less rack mount kit – Hot swappable power supplies and FAN trays – Power supply redundancy – Remote power cycling capability	– Throughput: 51.2 Tbps unidirectional – Buffer: 165 MB – Tool less rack mount kit – Hot swappable power supplies and FAN trays – Power supply redundancy – Remote power cycling capability

Features and benefits

AI/ML design

Artificial intelligence puts new challenges on compute, network, and storage solutions with large models that run in parallel across many GPUs for training. These models require fast job completion time (JCT) with minimal delays for the last GPU to finish its calculations, that is, low tail latency. Architects optimize the cluster performance through rail-optimized design ([Read this HPE Juniper Networking white paper](#) for more information about AI/ML cluster design). As model sizes and datasets continue to grow, designs must accommodate more GPUs in the cluster, requiring that the network seamlessly scale, without compromising performance, or introducing communication bottlenecks.

The QFX5240 line meets the needs of these large-scale AI networks. The switch provides:

- 64 ports of 800GbE on a 2U switch to reduce costs on both space and total power utilization
- Choice of connectivity with both OSFP and QSFP-DD variants of 800GbE for leaf-spine connectivity
- Advanced telemetry capabilities to support ECN/PFC counters
- Fine-grained, load-balancing capability to handle reduced flow entropy
- Automation of rail-optimized design through Apstra Data Center Director

Automation Data Center Director

Automation tools, such as Data Center Director, ensure the reliable setup of expansive networks with ongoing verification of the deployment along with monitoring of operations. Data Center Director delivers full day 0 through day 2+ capabilities for IP/EVPN fabrics with closed-loop assurance in the data center. Data Center Director provides a broad set of operational capabilities, with multiple built-in intent-based analytics probes, flow visibility, and analysis to ensure that the AI network is running as designed. Data Center Director provides a simple UI workflow to create custom intent-based analytics to capture, enrich, and visualize data from the AI network.

Monitoring

The QFX5240 line of switches supports [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

Additionally, the [Junos OS Evolved](#) supports a robust API set to support automation through Terraform, Ansible, zero-touch provisioning (ZTP), operations and event scripts, automatic rollback, and Python scripts.



QFX5240-64OD line of switches



QFX5240-64QD line of switches



QFX5241-32OD line of switches

* Product configuration and appearance may vary

Specifications

Hardware specifications

Table 2. QFX5240 line of switches system capacity

Parameter	QFX5240-64OD QFX5241-64OD	QFX5240-64QD QFX5241-64QD	QFX5241-32OD
System throughput	51.2/102.4 Tbps unidirectional/bidirectional	51.2/102.4 Tbps unidirectional/bidirectional	25.6/51.2 Tbps unidirectional/bidirectional
Max Forwarding rate	21.2 Bpps	21.2 Bpps	10.6 Bpps
Port density	64 OSFP 800GbE + 2 SFP+ 10GbE	64 QSFP 800GbE + 2 SFP+ 10GbE	32 ports of OSFP 800GbE + 2 SFP+ 10GbE
Max ports with breakout	64 × 800GbE, 128 × 400GbE, 256 × 200GbE, 256 × 100GbE. Max port density with mixed speeds 320 (please use online port checker to validate mixed speed configurations)	64 × 800GbE, 128 × 400GbE, 256 × 200GbE, 320 × 100GbE, 256 × 50GbE. Max port density with mixed speeds 320 (please use online port checker to validate mixed speed configurations)	32 × 800GbE, 64 × 400GbE, 128 × 200GbE, 160 × 100GbE, 128 × 50GbE. Max port density with mixed speeds 160 (please use online port checker to validate mixed speed configurations)
Dimensions (W x H x D)	17.26 x 3.46 x 25.52 in. (43.8 x 8.8 x 64.8 cm)	17.26 x 3.46 x 25.52 in. (43.8 x 8.8 x 64.8 cm)	17.25 × 1.69 × 23.19 in. (43.4 × 4.3 × 58.9 cm)
Rack units	2U	2U	1U
Weight	22 kg (48.50 lb) fully loaded without optics	22 kg (48.50 lb) fully loaded without optics	15 kg (33.17 lb)
Operating system	Junos OS Evolved	Junos OS Evolved	Junos OS Evolved
Switch chip	Broadcom Tomahawk 5	Broadcom Tomahawk 5	Broadcom Tomahawk 5
CPU	Intel® Ice Lake (4 core)	Intel Ice Lake (4 core)	Intel Ice Lake (4 core)
Memory	32 GB (16 GB x 2) of DDR4	32 GB (16 GB x 2) of DDR4	32 GB (16 GB x 2) of DDR4
Storage	2 x 480 GB	2 x 480 GB	2 x 480 GB
Power	Redundant (1+1) hot-pluggable 3000W AC 180 VAC—264 VAC single phase AC power	Redundant (1+1) hot-pluggable 3000W AC (200 to 240V) power supplies 180 VAC—264 VAC single phase AC power	Redundant (1+1) hot-pluggable 2400W power supplies 200 VAC—240 VAC single phase AC power
Cooling	Ports-to-PSU (AFO) 4 hot- pluggable fan modules	Ports-to-PSU (AFO) 4 hot- pluggable fan modules	Ports-to-PSU (AFO) 7 hot-pluggable fan modules
Total packet buffer	165 MB	165 MB	83 MB
Warranty	HPE standard one-year warranty	HPE standard one-year warranty	HPE standard one-year warranty

Table 3. QFX5240 line of switches feature matrix

Feature
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
VLAN—IEEE 802.1Q VLAN trunking
Routed VLAN interface (RVI)
Static MAC address assignment for interface
Global MAC learning disable
Link Aggregation and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad)
IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
Layer 3 features
Static routing
OSPF v2/v3
Filter-based forwarding
VRRP/VRRPv3
IPv6
Virtual routers
Loop-free alternate (LFA)
BGP
IS-IS
Dynamic Host Configuration Protocol (DHCP) v4/v6 relay (stateless)
VRF-aware DHCP
IPv4/IPv6 over GRE tunnels
Multicast
Internet Group Management Protocol (IGMP) v1/v2/v3
Multicast Listener Discovery (MLD) v2
IGMP proxy, querier
IGMP v1/v2/v3 snooping
Intersubnet multicast using IRB interface
MLD snooping
Protocol Independent Multicast PIM-SM, PIM-SSM, PIM-DM, PIM-Bidir Multicast Source Discovery Protocol (MSDP)
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

Table 3. QFX5240 line of switches feature matrix (continued)

Feature
10 hardware queues per port (8 unicast and 2 multicast)
Strict priority queuing (LLQ), shaped-deficit weighted round robin (SDWRR)
Layer 2 classification criteria: Interface, MAC address, Ether type, 802.1p, VLAN
Congestion avoidance capabilities: WRED, ECN
Trust IEEE 802.1p
Configurable shared buffer and buffer monitoring
Congestion Notification Profile
Priority-based flow control (PFC)—IEEE 802.1Qbb
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
Junos OS Evolved configuration rescue and rollback
Image rollback
SNMP v1/v2/v3
Junos OS Evolved XML management protocol
Automation and orchestration
Secure zero-touch provisioning (sZTP) (DevID)
Secure-BIOS
Secure-boot
Python
Junos OS Evolved event, commit, and OP scripts
Network services
ROCEv2
DCQCN, PFC, ECN
PFC watchdog
Dynamic load balancing (DLB)
Configurable hash-bucket size

Table 3. QFX5240 line of switches feature matrix (continued)

Feature
Software
– Operating System: Junos OS Evolved (recommended releases)
– Latency: 752 ns cut-through and 949 ns (Store and Forward)
Mode
– MAC addresses per system: 136K
– VLAN IDs: 4,096
– Number of link aggregation groups (LAGs): 128
– Number of ports per LAG: 64
– Firewall filters: 9,000
• Ingress: 18K Routed ACL (RACL), 2,250 VLAN ACL (VACL) and 18K Port ACL (PACL) rules
• Egress: 4K Routed ACL (RACL), 512 VLAN ACL (VACL) and 4K Port ACL (PACL) rules
– Address Resolution Protocol (ARP) entries: 32,000
– Generic routing encapsulation (GRE) tunnels: 1,000
– Jumbo frame: 9216 bytes
– Traffic mirroring
• Mirroring destination ports per switch: 3
• Maximum number of mirroring sessions: 7
• Mirroring destination VLANs per switch: 3
Note: Secure boot is available on QFX5241 switches

Environmental ranges

Table 4. QFX5240 line of switches operating parameters

Operating temperature	0° to 40°C @ sea level
Storage temperature	–40° to 70°C
Operating altitude	Up to 6000 ft/3962m
Relative humidity operating	5 to 90% non-condensing
Relative humidity nonoperating	5 to 90% non-condensing
Seismic	Zone 4 earthquake rating (GR-63 EQ zone 4)
QFX5240-64OD/QFX5241-64OD	
Typical power consumption	613W
Maximum power consumption	2116W
QFX5240-64QD/QFX5241-64QD	
Typical power consumption	616W
Maximum power consumption	2242W
QFX5241-32OD	
Typical power consumption	537W
Maximum power consumption	1381W

Notes:

Max power consumption measured at 40°C ambient temperature with DR optics at 100% load with IMIX traffic.
 Typical power consumption measured at 25°C ambient temperature with DACs at 50% load with IMIX traffic, excluding transceivers.
 Power consumption is subject to operating condition and unit-to-unit variations.

Electromagnetic compatibility

- FCC 47 CFR Part 15
- ICES-003 / ICES-GEN
- BS EN 55032
- BS EN 55035
- EN 300 386 V1.6.1
- EN 300 386 V2.2.1
- 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

Environmental compliance

- Restriction of Hazardous Substances (RoHS)
- Toxic Substances Control Act (TSCA)
- Persistent Organic Pollutants (POPs)
- Recycled Material Waste Electronics and Electrical Equipment (WEEE)
- Registration, Evaluation, Authorization and Restriction of Chemicals (REACH)
- Substances of Concern in Products (SCIP)

Safety

- UL 60950-1:2007 R10.14 Information Technology Equipment
- CAN/CSA-C22.2 No. 60950-1-07, Amd 1:2011, Amd 2:2014 Information Technology Equipment
- IEC 62368-1:2014 (2nd Edition) Audio/Video, Information and Communication Technology Equipment (Include all country deviation)
- IEC 62368-1:2018 (3rd Edition) Audio/Video, Information and Communication Technology Equipment (Include all country deviation)
- EN 62368-1:2014+A11:2017 Audio/Video Information and Communication Technology Equipment
- UL/CSA 62368-1:2019 (3rd Edition) Audio/Video, Information and Communication Technology Equipment
- IEC/EN 60825-1 Safety of Laser Products—Part 1: Equipment classification and requirements

ETSI

- ETSI 300753 (1997)—Acoustic noise emitted by telecommunications equipment

Ordering information

Product number	Description
QFX5241-32OD-AO	32x800GbE OSFP W/ SECURE BOOT, AC, AFO
QFX5241-32OD-DO	32x800GbE OSFP W/ SECURE BOOT, DC, AFO
QFX5241-32OD-DO-T1	32x800GbE OSFP, DC, Air-Out, ORv3
QFX5241-32OD-CHAS	32x800GbE W/ SECURE BOOT, CHASSIS Spare
QFX5241-32OD-TCHAS	32x800GbE OSFP, ORv3, Chassis Spare
QFX5241-64OD-DO-T2	64x800GbE OSFP, DC, Air-Out, ORv3
QFX5241-64OD-TCHAS	64x800GbE OSFP, ORv3 spare chassis without PSU, FAN
QFX5241-64QD-DO-T2	64x800GbE QSFP-DD, DC, Air-Out, ORv3
QFX5241-64QD-TCHAS	64x800GbE QSFP-DD, ORv3 spare chassis without PSU, FAN
QFX5241-1U-PWR-AO	QFX5241-1RU AC PSU, Air-Out
QFX5241-1U-PWR-DO	QFX5241-1RU DC PSU, Air-Out
QFX5241-1U-FAN-O	AIR OUT FANS FOR QFX5241-1RU
QFX5241-1U-4PRMK	4post Tool-less RMK for 1RU QFX5241
QFX5241-64OD-AO	64x800GbE OSFP, AC, AFO
QFX5241-64OD-DO	64x800GbE OSFP, DC, Air-Out
QFX5241-64QD-AO	64x800GbE QSFP-DD, AC, AFO
QFX5241-64QD-DO	64x800GbE QSFP-DD, DC, Air-Out
QFX5240-2U-4PRMK	4post Tool-less RMK for 2RU QFX5240 Series
QFX5240-2U-FANAO	Air Out FANS for QFX5240 Series-2RU
QFX5240-PWR-AC-AO	QFX5240 Series-2RU AC PSU, Air Out
QFX5241-PWR-DC-AO	QFX5240 Series-2RU DC PSU, Air Out
QFX5240-64OD-AO	64x800GE OSFP800, AC, Front-to-back airflow
QFX5241-64OD-CHAS	64x800GE OSFP800, spare chassis
QFX5240-64QD-AO	64x800GE QSFP-DD 800, AC, Front-to-back airflow
QFX5241-64QD-CHAS	64x800GE QSFP-DD 800, spare chassis
QFX5240-64OD-CHAS	64x800GE QSFP-DD 800, spare chassis
QFX5240-64QD-CHAS	64x800GE QSFP-DD 800, spare chassis
QFX5241-1U-FAN-AO	Air Out FANS for QFX5241-1RU

License SKUs

Product number	Description
S-QFX5K-C5-A1-X (X=3,5,P)	Advanced 1 Software License (X Years Subscription, X=3,5, or P for Perpetual) for QFX5240-OD/QD line of switches
S-QFX5K-C5-A2-X (X=3,5,P)	Advanced 2 Software License (X Years Subscription, X=1,3,5, or P for Perpetual) for QFX5240-OD/QD line of switches
S-QFX5K-C5-P1-X (X=3,5,P)	Premium Software License (X Years Subscription, X=1,3,5, or P for Perpetual) for QFX5240-OD/QD line of switches
S-QFX5K-C4-A1-X	SW, QFX5K for Advance 1 License for Host based overlays, Class 4 with SVC Customer support (X Years Subscription, X=3,5) for QFX5241-32OD line of switches
S-QFX5K-C4-A1-P	SW, QFX5K, Class 4 Products, Advance 1, Perpetual License, without Customer Support, must purchase CS SKU separately, Perpetual
S-QFX5K-C4-A2-X	SW, QFX5K for Advance 2 License for Network based overlays, Class 4 with SVC Customer support (X Years Subscription, X=3,5) for QFX5241-32OD line of switches
S-QFX5K-C4-A2-P	SW, QFX5K, Class 4 Products, Advance 2 license, without Customer Support, must purchase CS SKU separately, Perpetual
S-QFX5K-C4-P1-X	SW, QFX5K Premium License, Class 4 with SVC Customer support, (X Years Subscription, X=3,5) for QFX5241-32OD line of switches
S-QFX5K-C4-P1-P	SW, QFX5K, Class 4 Products, Premium 1 license, without Customer Support, must purchase CS SKU separately, Perpetual

Note: The information provided is from early prototyping and may vary from the actual GA product.

Optics and transceivers

Up-to-date information on supported optics can be found on the Hardware Compatibility Tool. The QFX5240 line of switches supports varying port speeds at 800GbE, 400GbE, and 100GbE, with different transceiver options.

Useful links

[Feature Explorer](#)

[Hardware Compatibility Tool](#)

[Recommended Releases](#)

[Port Checker](#)

[Power Calculator](#)

About HPE

HPE is a leader in essential enterprise technology, bringing together the power of AI, cloud, and networking to help organizations achieve more. As pioneers of possibility, our innovation and expertise advance the way people live and work. We empower our customers across industries to optimize operational performance, transform data into foresight, and maximize their impact. Unlock your boldest ambitions, with HPE. Discover more at [HPE.com](https://hpe.com).

[Visit HPE.com](https://hpe.com)

[Chat now](#)

© Copyright 2026 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel is a trademark of Intel Corporation or its subsidiaries in the U.S. and/or other countries. sFlow is a registered trademark of InMon Corp. All third-party marks are property of their respective owners.

a00150802ENW, Rev. 2

HEWLETT PACKARD ENTERPRISE

hpe.com

