ThinkSystem NVIDIA ConnectX-9 9180 800Gb XDR/2x400GbE OSFP PCIe Gen6 x16 Adapter

Product Guide

The ThinkSystem NVIDIA ConnectX-9 9180 800Gb XDR/2x400GbE OSFP PCIe Gen6 x16 Adapter is optimized to supercharge hyperscale AI computing workloads. With support for both InfiniBand and Ethernet networking at up to 800 Gb/s, ConnectX-9 SuperNIC delivers high-speed, efficient network connectivity, significantly enhancing system performance for AI factories and cloud data center environments. The ConnectX-9 9180 adapter is available as a PCIe adapter or OCP adapter (PCIe adapter planned for January 2027).

The following figure shows the ConnectX-9 9180 adapter in PCIe form factor (shown with the full-adapter gold case and internal heatsink removed)

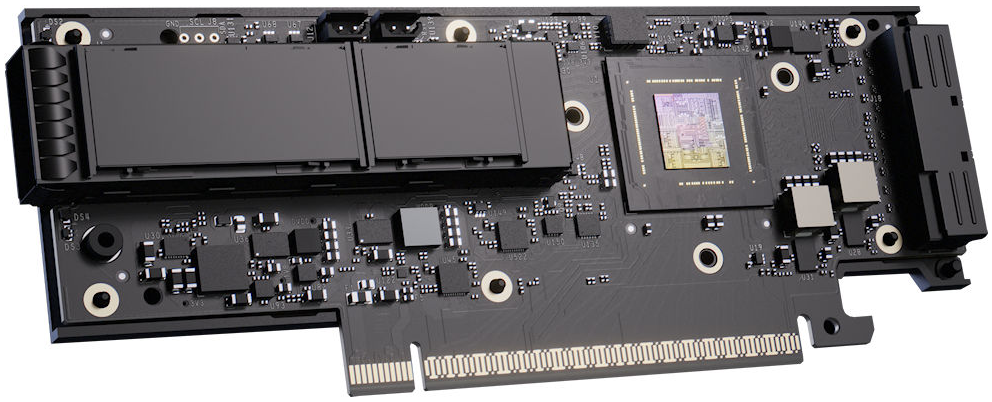


Figure 1. ThinkSystem NVIDIA ConnectX-9 9180 800Gb XDR/2x400GbE OSFP PCIe Gen6 x16 Adapter (full-adapter gold case and heatsink removed)

Did you know?

The ConnectX-9 9180 adapters are post-quantum cryptography (PQC) ready, providing security capabilities designed for the transition to cryptographic algorithms that can resist attacks from future quantum computers. The adapters support CNSA 2.0 post-quantum cryptography for secure boot and firmware update, helping protect the platform as security requirements evolve.

ConnectX-9 9180 provides up to 800 Gb/s connectivity using either InfiniBand or Ethernet, allowing the same adapter architecture to address different high-performance network infrastructures. With support for XDR InfiniBand and 800 GbE through a single OSFP connection, the ConnectX-9 9180 is well suited for the extremely high network bandwidth requirements of large-scale AI, HPC, and accelerated computing environments.

Part number information

The following table shows the ordering information.

Table 1. Ordering information

Part number	Feature code	Description	NVIDIA equivalent
Water-cooled adapter			
CTO only	CJP8	ThinkSystem NVIDIA ConnectX-9 9180 800Gbs XDR IB / 2x400GbE OSFP 1-Port OCP Gen6 x16 DWC	900-9X95E-00QB-MP0

The option part number includes the following:

- One NVIDIA adapter with full-height (3U) adapter bracket attached
- Low-profile (2U) adapter bracket
- Documentation

Supported transceivers and cables

The ConnectX-9 9180 adapter have an empty OSFP112 cage for connectivity.

The following table lists the supported transceivers.

Table 2. Transceivers

Part number	Feature code	Description	NVIDIA part number
400Gb Transceivers			
4TC7B10048	CFUE	Lenovo NVIDIA Single Port 400GbE OSFP, MPO12 APC, MMF Transceiver Up to 50m - Flat Top (Generic)	980-9I51S-F4NS00
4TC7A81826	BQJT	ThinkSystem NDR OSFP400 IB Multi Mode Solo-Transceiver	980-9I51S-00NS00
800Gb Transceivers with 2x 400Gb connections			
4TC7A83365	BQMJ	ThinkSystem NDRx2 OSFP800 IB Multi Mode Twin-Transceiver Flat Top	980-9I51A-00NS00
800Gb Transceivers			
4TC7B17648	CKW2	NVIDIA single port transceiver, DR4, 800Gbps, OSFP, MPO12(APC), 1310nm, up to 500m, RHS, TRO, Eth	980-9IAK0-00XM00
4TC7B10008	CCVE	Lenovo NVIDIA 800G XDR OSFP Single Mode Solo-Transceiver RHS (Generic)	980-9IAT0-00XM00

The following table lists the supported optical cables.

Table 3. Optical cables

Part number	Feature code	Description	NVIDIA part number
NVIDIA NDR Multi Mode Fibre Optical Cables			
4X97A81748	BQJN	Lenovo 3m NVIDIA NDR Multi Mode Optical Cable	980-9I73U-000003
4X97A81749	BQJP	Lenovo 5m NVIDIA NDR Multi Mode Optical Cable	980-9I73V-000005

Part number	Feature code	Description	NVIDIA part number
4X97A81750	BQJQ	Lenovo 7m NVIDIA NDR Multi Mode MPO12 APC Optical Cable	980-9I57W-000007
4X97A81751	BQJR	Lenovo 10m NVIDIA NDR Multi Mode Optical Cable	980-9I57X-00N010
4X97A81752	BQJS	Lenovo 20m NVIDIA NDR Multi Mode Optical Cable	980-9I57Z-000020
4X97A85349	BSN6	Lenovo 30M NVIDIA NDR Multi Mode MPO12 APC Optical Cable	980-9I570-00N030
NVIDIA XDR SM MPO12 APC Optical Cable			
4X97B11375	CCVD	Lenovo 3M NVIDIA SM MPO12 APC Optical Cable (Generic)	980-9I55A-00N003
4X97B05747	C98S	Lenovo 5M NVIDIA XDR SM MPO12 APC Optical Cable (Generic)	980-9I55B-00N005
4X97B05748	C98T	Lenovo 10M NVIDIA XDR SM MPO12 APC Optical Cable (Generic)	980-9I58D-00N010
4X97B05750	C98U	Lenovo 15M NVIDIA XDR SM MPO12 APC Optical Cable (Generic)	980-9I58E-00N015
4X97B05751	C98V	Lenovo 20M NVIDIA XDR SM MPO12 APC Optical Cable (Generic)	980-9I58F-00N020
4X97B05752	C98W	Lenovo 30M NVIDIA XDR SM MPO12 APC Optical Cable (Generic)	980-9I58G-00N030
4X97B09056	CB2T	Lenovo 50M NVIDIA SM MPO12 APC Optical Cable (Generic)	980-9I58H-00N050

The following table lists the supported direct-attach copper (DAC) cables.

Table 4. Copper cables

Part number	Feature code	Description	NVIDIA part number
NVIDIA NDRx2 OSFP800 Finned to NDRx2 OSFP800 Flat Copper Cables			
4X97A84581	BRKC	Lenovo 1M NVIDIA NDRx2 OSFP800 Finned to NDRx2 OSFP800 Flat Top Passive Copper Cable	980-9IA0H-00N001
4X97A84582	BRKD	Lenovo 1.5M NVIDIA NDRx2 OSFP800 Finned to NDRx2 OSFP800 Flat Top Passive Copper Cable	980-9IA0H-27N01A
4X97A84584	BRKF	Lenovo 3M NVIDIA NDRx2 OSFP800 Finned to NDRx2 OSFP800 Flat Top Active Copper Cable	980-9I601-00N003
NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Splitter Copper Cables			
4X97A81827	BQJV	Lenovo 1m NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	980-9I432-00N001
4X97A81828	BQJW	Lenovo 1.5m NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	980-9I926-00N01A
4X97A81829	BQJX	Lenovo 2m NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	980-9I924-00N002
4X97A81830	BQJY	Lenovo 3m NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	980-9I92N-00N003
NVIDIA XDR OSFP to 2xOSFP Active Copper Splitter Cable			
4X97B05741	C98M	Lenovo 1m NVIDIA XDR OSFP to 2xOSFP Active Copper Splitter cable (Generic)	980-9IAO5-00X001
4X97B05742	C98N	Lenovo 1.5m NVIDIA XDR OSFP to 2xOSFP Active Copper Splitter cable (Generic)	980-9IAO5-00X01A
4X97B05743	C98P	Lenovo 2m NVIDIA XDR OSFP to 2xOSFP Active Copper Splitter cable (Generic)	980-9IAO5-00X002
NVIDIA Twin Port XDR OSFP224 to 4x XDR400 OSFP RHS to RHS Linear Active Copper Splitter Cable			
4X97B17652	CLMF	Linear Active Copper Cables LACC Copper Cable 1,600G-to-4x400G, RHS OSFP-to- RHS OSFP, Switch to CX-9, 1m	980-9IAQ6-00X001
4X97B17653	CLMG	Linear Active Copper Cables LACC Copper Cable 1,600G-to-4x400G, RHS OSFP-to- RHS OSFP, Switch to CX-9, 2m	980-9IAQ6-00X002
NVIDIA Twin Port XDR OSFP224 to 2x XDR OSFP RHS to RHS Linear Active Copper Splitter Cable			
4X97B17654	CLMH	Linear Active Copper Cables LACC Copper Cable 1,600G-to-2x800G, RHS OSFP-to- RHS OSFP, Switch to CX-9, 1m	980-9IAO6-00X001
4X97B17655	CLMJ	Linear Active Copper Cables LACC Copper Cable 1,600G-to-2x800G, RHS OSFP-to- RHS OSFP, Switch to CX-9, 2m	980-9IAO6-00X002
NVIDIA Twin Port XDR OSFP224 to 4x XDR400 OSFP IHS to RHS Active Copper Splitter Cable			
4X97B17656	CLMK	Active Copper Splitter, 1m, OSFP to 4xOSFP, IHS to RHS	980-9IAQ5-00X001
4X97B17657	CLML	Active Copper Splitter, 2m, OSFP to 4xOSFP, IHS to RHS	980-9IAQ5-00X002

Features

The NVIDIA ConnectX-9 9180 adapters have the following features:

- **Up to 800 Gb/s network connectivity:** Provides a single OSFP network port supporting up to 800 Gb/s bandwidth. The 9180 adapters support both InfiniBand and Ethernet, enabling very high-bandwidth connectivity for AI, HPC, and cloud data center workloads.
- **InfiniBand and Ethernet support:** Supports XDR, NDR, NDR200, HDR, and HDR100 InfiniBand connectivity, as well as Ethernet operating at 800, 400, 200, and 100 Gb/s.
- **PCIe 6.0 host interface:** Uses a PCIe Gen6 x16 host interface, providing high-bandwidth connectivity between the adapter and host system. The interface operates at up to 64 GT/s per lane and is also compatible with PCIe Gen5.
- **Designed for accelerated AI networking:** Includes hardware capabilities for RDMA and RoCEv2 acceleration, programmable RDMA transport, and advanced programmable congestion control, providing efficient network connectivity for demanding AI fabrics and accelerated computing environments.
- **NVIDIA GPUDirect acceleration:** Supports NVIDIA GPUDirect RDMA and GPUDirect Storage, providing efficient data paths involving GPUs, network adapters, and storage. These capabilities are particularly applicable to GPU-intensive AI and HPC environments.
- **Advanced congestion control:** Provides advanced programmable and telemetry-based congestion-control capabilities designed to maintain network efficiency for large-scale AI workloads.
- **In-Network Computing acceleration:** Extends NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP) capabilities for In-Network Computing, helping accelerate supported collective operations for scientific computing, large-scale AI training, and inference workloads.
- **NVIDIA Spectrum-X integration:** Supports NVIDIA Spectrum-X Ethernet technologies, including Spectrum-X multi-plane and the Spectrum-X net plug-in, for integration into large-scale accelerated computing environments.
- **Virtualization and cloud networking acceleration:** Provides SR-IOV and NVIDIA Accelerated Switching and Packet Processing (ASAP^2) capabilities for software-defined networking and virtual network functions. Acceleration capabilities include Open vSwitch, overlay networking, connection tracking, and network address translation (NAT).
- **TCP offload capabilities:** Provides stateless TCP offloads to reduce host processing associated with networking functions and improve efficiency in cloud and data center networking environments.
- **Multi-Host connectivity:** Supports NVIDIA Multi-Host technology for connectivity to as many as four hosts, enabling flexible connectivity configurations in supported systems.
- **Socket Direct capability:** Supports NVIDIA Socket Direct technology, which can provide direct PCIe connectivity from separate CPUs in a dual-socket server to the network adapter, helping optimize network access in multi-socket server configurations.
- **Hardware-accelerated cryptography:** Provides AES-GCM acceleration using 128-bit and 256-bit keys for data-in-motion security, including IPsec, TLS, and PSP, together with AES-XTS 256/512-bit key acceleration for data-at-rest protection.
- **Advanced platform security:** Includes secure boot with a hardware root of trust, secure firmware update, flash encryption, and device attestation using SPDM 1.1 and SPDM 1.2. The 9180 adapters covered by this product guide are crypto enabled, secure boot enabled, and PQC ready.

- **Post-quantum cryptography (PQC) readiness:** Supports CNSA 2.0 post-quantum cryptography capabilities for secure boot and firmware update. The 9180 adapters covered by this product guide are PQC ready.
- **Advanced timing and synchronization:** Provides precision network timing and synchronization capabilities, including IEEE 1588 PTP, Precision Time Measurement (PTM), TCXO, PPS input/output, and SyncE.

Technical specifications

The ThinkSystem NVIDIA ConnectX-9 9180 800Gb XDR/2x400GbE OSFP PCIe Gen6 x16 Adapter have the following technical specifications:

Network Interface

- One OSFP112 cage for a transceiver or cable
- Max. bandwidth: 800 Gb/s
- 200/100/50 Gb/s of PAM4 and 25/10 Gb/s NRZ

Protocol support

- InfiniBand protocol support, based on InfiniBand Architecture Specification (IBTA) v1.7:
  - 800 Gb/s XDR, (4 lanes × 200 Gb/s per lane)
  - 400 Gb/s NDR (4 lanes × 100 Gb/s per lane)
  - 200 Gb/s NDR200 (2 lanes × 100 Gb/s per lane)
  - 100 Gb/s HDR100 (2 lanes × 50 Gb/s per lane)
  - 50 Gb/s HDR (50 Gb/s per lane)
- Ethernet protocols supported are as follows:
  - 800 Gb/s: 800GAUI-4 C2M, 800GAUI-8 C2M, 800GBASE-CR4, 800GBASE-CR8
  - 400 Gb/s: 400GAUI-2 C2M, 400GAUI-4 C2M, 400GAUI-8 C2M, 400GBASE-CR2, 400GBASE-CR4, 400GBASE-CR8
  - 200 Gb/s: 200GAUI-1 C2M, 200GAUI-2 C2M, 200GAUI-4, 200GAUI-4 C2M, 200GBASE-CR1, 200GBASE-CR2, 200GBASE-CR4
  - 100 Gb/s: 100GAUI-1 C2M, 100GAUI-2 C2M, CAUI-4 C2M, 100GBASE-CR1, 100GBASE-CR2, 100GBASE-CR4
  - 50 Gb/s: 50GAUI-1 C2M, 50GBASE-CR
  - 25 Gb/s: 25GAUI C2M, 25GBASE-CR, 25GBASE-CR-S

Host Interface

- PCIe Gen6 (up to 48 lanes)
- NVIDIA Multi-Host (up to 4 hosts)
- PCIe switch downstream port containment (DPC)
- Message-signaled interrupts (MSI) and MSI-X

AI Networking

- Remote direct-memory access (RDMA) and RDMA over Converged Ethernet v2 (RoCEv2) accelerations
- NVIDIA Spectrum-X Ethernet:
  - Spectrum-X multi-plane
  - Spectrum-X net plug-in
- Programmable RDMA transport
- Advanced, programmable congestion control

- NVIDIA Collective Communications Library (NCCL)
- NVIDIA Inference Transfer Library (NIXL)
- NVIDIA GPUDirect RDMA
- GPUDirect Storage
- In-Network Computing
- Advanced timing and synchronization
- Message Passing Interface (MPI) accelerations

Cloud Networking

- Stateless Transmission Control Protocol (TCP) offloads
- Single-root IO virtualization (SR-IOV)
- NVIDIA Accelerated Switching and Packet Processing (ASAP2) Ethernet for software-defined networking (SDN) and virtual network functions (VNF):
    - Open vSwitch (OVS) acceleration
    - Overlay network accelerations
    - Connection tracking (L4 firewall) and network address translation (NAT)

Security

- AES-GCM 128/256-bit key data-in-motion accelerations: Internet Protocol Security (IPsec), Transport Layer Security (TLS), Protocol for Secure Packet (PSP)
- AES-XTS 256/512-bit key data-at-rest acceleration
- Platform security:
    - Secure boot with hardware root of trust
    - Secure firmware update
    - Flash encryption
    - Device attestation (SPDM 1.1 and SPDM 1.2)
    - Commercial National Security Algorithm Suite (CNSA) 2.0 post-quantum cryptography (PQC) Secure Boot and Update support

Management and Control

- Network Control Sideband Interface (NC-SI) (OCP adapter only)
- Management Component Transport Protocol (MCTP) over System Management Bus (SMBus) and PCIe Platform Level Data Model (PLDM) for:
    - Monitor and Control DSP0248
    - Firmware Update DSP0267
    - Redfish Device Enablement DSP0218
    - Field-Replaceable Unit (FRU) DSP0257
- Security Protocols and Data Models (SPDM) DSP0274
- Serial Peripheral Interface (SPI) to flash

Network Boot

- InfiniBand or Ethernet
- Preboot eXecution Environment (PXE) boot
- Unified Extensible Firmware Interface (UEFI)

NVIDIA Unified Fabric Manager

NVIDIA Unified Fabric Manager (UFM) is InfiniBand networking management software that combines enhanced, real-time network telemetry with fabric visibility and control to support scale-out InfiniBand data centers. The two UFM offerings are available from Lenovo:

- **UFM Telemetry** for Real-Time Monitoring

  The UFM Telemetry platform provides network validation tools to monitor network performance and conditions, capturing and streaming rich real-time network telemetry information, application workload usage, and system configuration to an on-premises or cloud-based database for further analysis.

- **UFM Enterprise** for Fabric Visibility and Control

  The UFM Enterprise platform combines the benefits of UFM Telemetry with enhanced network monitoring and management. It performs automated network discovery and provisioning, traffic monitoring, and congestion discovery. It also enables job schedule provisioning and integrates with industry-leading job schedulers and cloud and cluster managers, including Slurm and Platform Load Sharing Facility (LSF).

The following table lists the UFM subscription licenses available from Lenovo.

Table 5. NVIDIA Unified Fabric Manager subscriptions

Part number	Feature code (7S09CTO6WW)	Description	NVIDIA part number
UFM Telemetry			
7S090011WW	S921	NVIDIA UFM Tel Subscription SW with Business Critical for up to 1000 InfiniBand Adapters (HCAs), 1 Year	797-YUFTMA+P3CMI12
7S090012WW	S922	NVIDIA UFM Tel Subscription SW with Business Critical for up to 1000 InfiniBand Adapters (HCAs), 3 Year	797-YUFTMA+P3CMI36
7S090013WW	S923	NVIDIA UFM Tel Subscription SW with Business Critical for up to 1000 InfiniBand Adapters (HCAs), 5 Year	797-YUFTMA+P3CMI60
UFM Enterprise			
7S09000XWW	S91Y	NVIDIA UFM Ent Subscription SW with Business Critical for up to 1000 InfiniBand Adapters (HCAs), 1 Year	797-YUFETA+P3CMI12
7S09000YWW	S91Z	NVIDIA UFM Ent Subscription SW with Business Critical for up to 1000 InfiniBand Adapters (HCAs), 3 Year	797-YUFETA+P3CMI36
7S09000ZWW	S920	NVIDIA UFM Ent Subscription SW with Business Critical for up to 1000 InfiniBand Adapters (HCAs), 5 Year	797-YUFETA+P3CMI60

For more information on NVIDIA UFM, see the NVIDIA Unified Fabric Manager product page.

Server support

Server support for the ConnectX-9 9180 adapter is as follows:

Table 6. Server support - Current servers (Part 1 of 2)

Part Number	Description	AMD V3					2S Intel V3/V4						1S V3				4S 8S Intel V3/V4				
		SR635 V3 (7D9H / 7D9G)	SR655 V3 (7D9F / 7D9E)	SR645 V3 (7D9D / 7D9C)	SR665 V3 (7D9B / 7D9A)	SD535 V3 (7DD8 / 7DD1)	ST650 V3 (7D7B / 7D7A)	SR630 V3 (7D72 / 7D73)	SR650 V3 (7D75 / 7D76)	SR630 V4 (7DG8 / 7DG9)	SR650 V4 (7DGC / 7DGD)	SR650a V4 (7DGC / 7DGD)	ST45 V3 (7DH4 / 7DH5)	ST50 V3 (7DF4 / 7DF3)	ST250 V3 (7DCF / 7DCE)	SR250 V3 (7DCM / 7DCL)	SR850 V3 (7D97 / 7D96)	SR860 V3 (7D94 / 7D93)	SR950 V3 (7DC5 / 7DC4)	SR850 V4 (7DJT / 7DJS)	SR860 V4 (7DJQ / 7DJN)

Table 7. Server support - Current servers (Part 2 of 2)

Part Number	Description	AI Servers					Edge		Super computing		
		SR675 V3 (7D9Q / 7D9R)	SR680a V3 B200 (7DM9)	SR685a V3 (7DHC)	SR780a V3 (7DJ5)	SR680a V4 (7DMK)	SE100 (7DGR)	SE455 V3 (7DBY)	SC750 V4 (7DDJ)	SC777 V4 (7DKA)	SD665 V3 (7D9P)

- ThinkSystem SC755 V5

Support of adapters with generic firmware

One or more of the adapters described in this product guide uses standard vendor firmware (look for "Generic FW" or "Generic" in the adapter names). These adapters are supported in Lenovo servers and with Lenovo management tools.

Support in Lenovo XClarity management tools for adapters with generic firmware is per the following table.

Tip: Always use firmware that is obtained from Lenovo sources to ensure the firmware is fully tested by Lenovo and is supported. You should not use firmware that is obtained from the vendor web site, unless directed to do so by Lenovo support.

Table 11. Lenovo XClarity management tools support for adapters with generic firmware

Function	Lenovo XClarity Provisioning Manager	Lenovo XClarity OneCLI (out-of-band)	Lenovo XClarity OneCLI (in-band)	Lenovo XClarity Administrator
Adapter configuration	Supported (in-band via UEFI)	Supported with latest UEFI firmware*	Supported with latest UEFI firmware*	Supported with latest UEFI firmware*

* Minimum supported UEFI firmware is as follows: SR630 V4, SR650 V4, SR650a V4: Release IHE116I; SR850 V4, SR860 V4: Release RVE106I.

Regulatory approvals

The adapter has the following regulatory approvals:

- Safety: CB / cTUVus / CE
- EMC: CE / FCC / VCCI / ICES / RCM / KC
- RoHS: RoHS Compliant

Operating environment

The adapter has the following operating characteristics:

- Maximum power available through OSFP port: 17W
- Temperature
  - Operational: 0°C to 55°C
  - Non-operational: -40°C to 70°C
- Humidity: 90% relative humidity

Warranty

One year limited warranty. When installed in a Lenovo server, the adapter assumes the server's base warranty and any warranty upgrades.

Related publications

For more information, refer to these documents:

- ThinkSystem Ethernet and InfiniBand Adapter Reference
  https://lenovopress.lenovo.com/lp1594-thinksystem-ethernet-infiniband-adapter-reference
- Lenovo ServerProven compatibility information
  http://serverproven.lenovo.com
- NVIDIA ConnectX InfiniBand product page:
  https://www.nvidia.com/en-us/networking/infiniband-adapters/
- ConnectX-9 SuperNIC user manual:
  https://networking-docs.nvidia.com/connectx9hw/
- ConnectX-9 SuperNIC for OCP 3.0 user manual:
  https://networking-docs.nvidia.com/connectx9ocphw/

Related product families

Product families related to this document are the following:

- InfiniBand & Omni-Path Adapters

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION ”AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2026. All rights reserved.

This document, LP2561, was created or updated on November 10, 2026.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.lenovo.com/LP2561
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.lenovo.com/LP2561.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
ServerProven®
ThinkSystem®
XClarity®

The following terms are trademarks of other companies:

AMD® is a trademark of Advanced Micro Devices, Inc.

Intel®, the Intel logo is a trademark of Intel Corporation or its subsidiaries.

NVIDIA®, NVIDIA Socket Direct, NVIDIA Spectrum®, and NVIDIA Spectrum-X are trademarks of NVIDIA Corporation.

Other company, product, or service names may be trademarks or service marks of others.