




ConnectX-7 400G Adapters

Accelerated networking for modern data center infrastructures.



The NVIDIA® ConnectX®-7 family of remote direct-memory access (RDMA) network adapters supports InfiniBand and Ethernet protocols and a range of speeds up to 400 gigabits per second (Gb/s). It enables a wide range of advanced, scalable, and secure networking solutions that address traditional enterprise needs up to the world's most demanding AI, scientific computing, and hyperscale cloud data center workloads.

Accelerated Networking and Security

ConnectX-7 provides a broad set of software-defined, hardware-accelerated networking, storage, and security capabilities, which enable organizations to modernize and secure their IT infrastructures. ConnectX-7 also powers agile and high-performance solutions from edge to core data centers to clouds, all while enhancing network security and reducing total cost of ownership.

Accelerate Data-Driven Scientific Computing

ConnectX-7 provides ultra-low latency, extreme throughput, and innovative NVIDIA In-Network Computing engines to deliver the acceleration, scalability, and feature-rich technology needed for today's modern scientific computing workloads.

Features*

InfiniBand Interface

- > InfiniBand Trade Association Spec 1.5 compliant
- > RDMA, send/receive semantics
- > 16 million input/output (IO) channels
- > 256 to 4 kilobyte maximum transmission unit (MTU), 2 Gb messages

Ethernet Interface

- > Up to four network ports supporting non-return-to-zero (NRZ) and four-level pulse amplitude-modulated signal (PAM4), both 50G and 100G, in various configurations
- > Up to 400 Gb/s total bandwidth
- > RDMA Over Converged Ethernet (RoCE)

Product Specifications

Supported Network Protocols	InfiniBand, Ethernet
Total Bandwidth	400 Gb/s
InfiniBand Speeds	NDR 400 Gb/s, HDR 200 Gb/s, EDR 100 Gb/s
Ethernet Speeds	400 GbE, 200 GbE, 100 GbE, 50 GbE, 25 GbE, 10 GbE
Number of Network Ports	1, 2, or 4
Host Interface	PCIe Gen5, up to 32 lanes
Form Factors	PCIe HHHL/FHHL and OCP3.0 TSFF/SFF
Interface Technologies	NRZ (10G, 25G), PAM4 (50G, 100G)

Enhanced InfiniBand Networking

- > Hardware-based reliable transport
- > Extended reliable connected (XRC) transport
- > Dynamically connected transport (DCT)
- > NVIDIA GPUDirect® RDMA
- > GPUDirect Storage
- > Adaptive routing support
- > Enhanced atomic operations
- > Advanced memory mapping, allowing user mode registration (UMR)
- > On-demand paging (ODP), including registration-free RDMA memory access
- > Enhanced congestion control
- > Burst buffer offload
- > Single-root IO virtualization (SR-IOV)
- > Optimized for high-performance computing (HPC) software libraries, including:
 - NVIDIA HPC-X®, NVIDIA Unified Communication X (UCX®), NVIDIA Unified Collective Communication (UCC), NVIDIA Collective Communications Library (NCCL), OpenMPI, MVAPICH, MPICH, OpenSHMEM, partitioned global address space (PGAS)
- > Collective operations offloads
- > Support for NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™
- > Rendezvous protocol offload
- > In-network on-board memory

Remote Boot

- > Remote boot over InfiniBand
- > Remote boot over Internet Small Computer Systems Interface (iSCSI)
- > Unified Extensible Firmware Interface (UEFI)
- > Preboot Execution Environment (PXE)

Enhanced Ethernet Networking

- > Zero-touch RoCE
- > NVIDIA Accelerated Switch and Packet Processing (ASAP²)™ for software-defined networking (SDN) and virtual network functions (VNF)
 - Open vSwitch (OVS) acceleration
 - Overlay network acceleration: Virtual Extensible LAN (VXLAN), Generic Network Virtualization Encapsulation (GENEVE), Network Virtualization Using Generic Routing Encapsulation (NVGRE)
 - Connection tracking (L4 firewall) and network address translation (NAT)
 - Flow mirroring, header rewrite, hierarchical quality of service (QoS)
- > SR-IOV
- > Stateless Transmission Control Protocol (TCP) offloads

Storage Accelerations

- > Block-level encryption: XEX-based tweaked codebook mode with ciphertext stealing-Advanced Encryption Standard (XTS-AES) 256/512-bit key
- > Non-Volatile Memory Express over Fabrics (NVMe-oF) and NVMe/TCP acceleration
- > T10 Data Integrity Field (T10-DIF) signature handover
- > Server Routing Protocol (SRP), iSCSI Extensions for RDMA (iSER), Network File System (NFS) over RDMA, Server Message Block (SMB) Direct

Management and Control

- > Network controller sideboard interface (NC-SI), Management Component Transport Protocol (MCTP) over System Management Bus (SMBus), and MCTP over PCIe
- > Platform Level Data Model (PLDM) for Monitor and Control DSP0248
- > PLDM for Firmware Update DSP0267
- > PLDM for Redfish Device Enablement DSP0218
- > PLDM for Field-Replaceable Unit (FRU) DSP0257
- > Security Protocols and Data Models (SPDM) DSP0274
- > Serial Peripheral Interface (SPI) to flash
- > Joint Test Action Group (JTAG) Institute of Electrical and Electronics Engineers (IEEE) 1149.1 and IEEE 1149.6

Cybersecurity

- > Platform security:
 - Secure boot with hardware root of trust
 - Secure firmware update
 - Flash encryption
 - Device attestation
- > Internet Protocol Security (IPsec)/Transport Layer Security (TLS)/Media Access Control Security (MACSec) 128/256-bit key data-in-motion encryption
- > IPsec for RoCE and Ethernet

Advanced Timing and Synchronization

- > Advanced Precision Time Protocol (PTP): IEEE 1588v2 (any profile), G.8273.2 Class C, line-rate hardware timestamp (UTC format)
- > Precise Time Measurement (PTM)
- > SyncE: Meets G.8262.1 (eEEEC)
- > Configurable pulse per second (PPS) in and out
- > Time-triggered scheduling
- > PTP-based packet pacing

Compatibility

PCI Express Interface

- > PCIe Gen 5.0 compatible, 32 lanes
- > Support for PCIe bifurcation
- > NVIDIA Multi-Host™ supports connection of up to four hosts
- > PCIe switch downstream port containment (DPC)
- > Support for message-signaled interrupts (MSI)/MSI-X mechanisms

Portfolio and Ordering Information

The portfolio of ConnectX-7 network adapters and ordering information is available in the ConnectX-7 user manuals: [PCIe adapters manual](#) and [OCP 3.0 adapters manual](#).

Ready to Get Started?

To learn more Email GOPNY@pny.com

© 2025 NVIDIA Corporation and affiliates. All rights reserved. NVIDIA, the NVIDIA logo, Accelerated Switch and Packet Processing (ASAP2), ConnectX, GPUDirect, HPC-X, Multi-Host, Scalable Hierarchical Aggregation and Reduction Protocol (SHARP), and UCX are trademarks and/or registered trademarks of NVIDIA Corporation. All company and product names are trademarks or registered trademarks of the respective owners with which they are associated. Features, pricing, availability, and specifications are all subject to change without notice. 4073150. JUL25

