Datasheet

ConnectX-7 2x200G Network Adapter

Accelerated networking for modern data center infrastructures.

NVIDIA® ConnectX®-7 with two 200Gb/s ports is a remote direct-memory access (RDMA) network adapter that supports Ethernet and InfiniBand protocols and a range of speeds up to 200Gb/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. Moreover, ConnectX-7 powers agile and high-performance solutions from edge to core data centers to clouds, all while enhancing network security and reducing the 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.

Product Specifications	
Supported network protocols	InfiniBand, Ethernet
Ethernet speeds*	200GbE, 100GbE, 50GbE, 25GbE, 10GbE
InfiniBand speeds	NDR200 200Gb/s, HDR 200Gb/s, EDR 100Gb/s
Number of network ports	2
Host interface	PCIe Gen 5, up to 32 lanes
Form factor	PCIe HHHL
Network interface technologies	NRZ (10G, 25G), PAM4 (50G, 100G)

*Lower speeds are enabled via QSFP to SFP (QSA) modules.

Features*

Ethernet Interface

- Two network ports supporting NRZ, PAM4 (50G and 100G)
- Up to 400Gb/s total bandwidth
- RDMA over converged Ethernet (RoCE)
- Ethernet mode is the default; for modification instructions refer to the **NVIDIA ConnectX-7 Adapter Cards User Manual.**

InfiniBand Interface

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

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)
  - Flow mirroring, header rewrite, hierarchical quality of service (QoS)
- Single-root IO virtualization (SR-IOV)
- Stateless Transmission Control Protocol (TCP) offloads

Enhanced InfiniBand Networking

- Hardware-based reliable transport
- Extended Reliable Connected (XRC)
- 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

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), Internet Small Computer Systems Interface (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

Remote Boot

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

Cybersecurity

- Platform security
  - Secure boot with hardware root of trust
  - Secure firmware update
  - Flash encryption
  - Device attestation

Compatibility

PCI Express Interface

- PCIe Gen 5 compatible, 32 lanes through extension card
- 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

Operating Systems/Distributions

- In-box drivers for major operating systems
  - Linux: RHEL, Ubuntu
  - Windows
- Virtualization and containers
  - VMware ESXi (SR-IOV)
  - Kubernetes

Ordering Information

For product specifications visit the online **ConnectX-7 PCIe user manual**.

Cisco SKU	Cisco PID	NVIDIA SKU	Product Description	Qualified Cisco Servers
30-100325-01	UCSC-P-N7D200GF	900-9X7AH-0078-DTZ (Legacy OPN: MCX755106AS-HEAT)	ConnectX-7 HHHL Adapter Card, 200GbE (default mode) / NDR200 IB, Dual-port QSFP112, PCIe 5 x16 with x16 PCIe extension option, Crypto Disabled, Secure Boot Enabled, Tall bracket	Cisco rack servers: UCS C220 M7, UCS C240 M7

Cards come assembled with a tall bracket PN MEC018771, if needed the short bracket PN MEC016919 is available in the box as well.
Server qualification is presented as of the date of publication. For latest server qualification information, please visit the **Cisco Hardware Compatibility List (HCL)**.

Optics and Cables

NVIDIA-Supported Cables and Modules

NVIDIA recommends using NVIDIA cables and modules. **Additional information on tested modules**:
Select firmware version > Select NVIDIA OPN > Select PSID > Select "Release Notes" under Download/Documentation.

Cisco-Branded Interoperable Cables and Modules: basic Interoperability by NVIDIA & UCS at NIC introduction.

QSFP Product ID	Product Description
QSFP-100G-AOC10M	100GBASE QSFP active optical cable, 10m
QSFP-100G-AOC7M	100GBASE QSFP active optical cable, 7m
QSFP-100G-AOC5M	100GBASE QSFP active optical cable, 5m
QSFP-100G-CU5M	100GBASE-CR4 passive copper cable, 5m
QSFP-100G-CU3M	100GBASE CR4 passive copper cable, 3m
QSFP-100G-LR4-S	100GBASE-LR4 QSFP transceiver, LC, 10km over SMF
QSFP-100G-SR4-S	100GBASE SR4 QSFP transceiver, MPO, 100m over OM4 MMF
QSFP-40/100-SRBD	100G and 40GBASE SR-BiDi QSFP transceiver, LC, 100m OM4 MMF
QSFP-100G-DR-S	100GBASE DR QSFP transceiver, 500m over SMF
QSFP-100G-LR4-S	100GBASE LR4 QSFP transceiver, LC, 10km over SMF
QSFP-100G-SR1.2	100GBASE QSFP, Bi-Direction Transceiver over MMF
QSFP-200-CU3M	200G QSFP56 to QSFP56, Passive Copper Cable 3m

For latest updates on the TMG Supported Optics and Cables, please consult the **Cisco Transceiver Module Group (TMG) Compatibility Matrix**.

Break-out and split cables are not supported.

Compatibility-Tested Network Switches

- Cisco Nexus 93180YC-FX Switch
- Cisco Nexus 93180YC-EX Switch
- Cisco Nexus C93600CD-GX Switch
- Cisco Nexus 5596UP Switch
- Cisco Nexus 5596T Switch

Tested 200GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
200GbE	Spectrum-3	MSN4600V-XXXX	64 QSFP56 ports, 200GbE 2U open Ethernet switch with Onyx	NVIDIA
200GbE	Spectrum-2	MSN3700-XXXX	32 QSFP56 ports, 200GbE open Ethernet switch system	NVIDIA

Tested 100GbE Switches

Speed	Switch Silicon	OPN # / Name	Description	Vendor
100GbE	Spectrum-3	MSN4600-XXXX	64-port non-blocking 100GbE open Ethernet switch system	NVIDIA
100GbE	Spectrum-2	MSN3700C-XXXX	32-port non-blocking 100GbE open Ethernet switch system	NVIDIA
100GbE	Spectrum-2	MSN3420-XXXX	48 SFP + 12 QSFP ports non-blocking 100GbE open Ethernet switch system	NVIDIA
100GbE	Spectrum	MSN2700-XXXX	32-port non-blocking 100GbE open Ethernet switch system	NVIDIA
100GbE	N/A	QFX5200-32C-32	32-port 100GbE Ethernet switch system	Juniper
100GbE	N/A	7060CX-32S	32-port 100GbE Ethernet switch system	Arista
100GbE	N/A	3232C	32-port 100GbE Ethernet switch system	Cisco
100GbE	N/A	N9K-C9236C	36-port 100GbE Ethernet switch system	Cisco
100GbE	N/A	93180YC-EX	48-port 25GbE + 6-port 100GbE Ethernet switch system	Cisco
100GbE	N/A	S6820-56HF	L3 Ethernet switch with 48 SFP28 ports and 8 QSFP28 ports	H3C
100GbE	N/A	BMS T7032-IX7	32 QSFP28 ports support for 10/25/40/50/100GbE	QuantaMesh

Environmental Specifications

Temperature	
Operational: 0°C to 55°C (32°F to 131°F)	**Non-operational:** –40°C to 70°C (–40°F to 158°F)

The non-operational storage temperature specifications apply to the product without its package.

Electrical and Thermal Specifications

Cable	
Typical power with passive cables in PCIe Gen 5 x16	24.9W

Power numbers are provided for passive cables only. For board power numbers while using active cables, please use the following formula:
TMP with Active Cables = TMP with Passive Cable + Active Module Power x Number_of_Modules / 0.9 (efficiency factor).

Ready to Get Started?

To learn more about InfiniBand adapters, visit:
nvidia.com/infiniband-adapters

To learn more about Ethernet network interface cards (NICs), visit:
nvidia.com/ethernet-adapters

* This section describes hardware features and capabilities.
Please refer to the **driver** and firmware release notes for feature availability.
Images are for illustration only; actual products may vary.