




NVIDIA CONNECTX-7

InfiniBand In-Network Computing Adapter



Accelerate Data-Driven Scientific Computing with In-Network Computing


As a key component of the NVIDIA® Quantum-2 InfiniBand platform, the NVIDIA ConnectX®-7 smart host channel adapter (HCA) provides the highest networking performance available to take on the world's most challenging problems. The ConnectX-7 InfiniBand adapter provides ultra-low latency, 400Gb/s throughput, and innovative NVIDIA In-Network Computing engines to deliver the acceleration, scalability, and feature-rich technology needed for high performance computing (HPC), artificial intelligence (AI), and hyperscale cloud data centers.

High performance computing and artificial intelligence have driven supercomputers into wide commercial use as the primary data processing engines enabling research, scientific discoveries and product development. These systems can carry complex simulations and unlock the new era of AI, where software writes software. NVIDIA InfiniBand networking is the engine of these platforms delivering breakthrough performance.

ConnectX-7's smart In-Network Computing acceleration engines include collective accelerations, MPI Tag Matching and MPI_Alltoall engines. These performance advantages and the standard guarantee of backward- and forward-compatibility ensure leading performance and scalability for compute and data-intensive applications and enable users to protect their data center investments.

Portfolio

- > Single-port or dual-port 400Gb/s or 200Gb/s, with octal small form-factor pluggable (OSFP) connectors
- > Dual-port 200Gb/s with quad small form-factor pluggable (QSFP) connectors
- > PCIe standup half-height, half-length (HHHL) form factor, with options for NVIDIA Socket Direct™
- > Open Compute Project 3.0 (OCP3.0) tall small form factor (TSFF) and small form factor (SFF)
- > Standalone ConnectX-7 application-specific integrated circuit (ASIC), supporting PCIe switch capabilities



PRODUCT SPECIFICATIONS

Max total bandwidth	400Gb/s
IBTA Spec compliant	1.5
Number of network ports	1/2/4
Host interface	PCIe Gen5, up to x32 lanes
RDMA message rate	330-370 million messages per second
Acceleration engines	Collective operations MPI_Alltoall MPI tag matching
Advanced storage capabilities	Block level encryption and checksum offloads
Accurate timing	PTP 1558v2, 16ns accuracy
Secure boot	On-chip with hardware root-of-trust
Host management	NC-SI, MCTP over SMBus, and MCTP over PCIe
Supported operating systems	Linux, Windows, VMware
Form factors	PCIe HHHL, Socket Direct OCP3.0 TSFF, SFF

Features*

InfiniBand

- > InfiniBand Trade Association (IBTA) Specification 1.5 compliant
- > Up to four ports
- > Remote direct-memory access (RDMA), send/receive semantics
- > Hardware-based congestion control
- > Atomic operations
- > 16 million input/output (IO) channels
- > 256 to 4Kbyte maximum transmission unit (MTU), 2Gbyte messages
- > 8x virtual lanes (VL) + VL15

Enhanced Networking

- > Hardware-based reliable transport
- > Extended Reliable Connected (XRC) transport
- > Dynamically Connected Transport (DCT)
- > GPUDirect® RDMA
- > GPUDirect Storage
- > Out-of-order RDMA supporting adaptive routing
- > Enhanced atomic operations
- > Advanced memory mapping support, allowing user mode registration and remapping of memory (UMR)
- > On-demand paging (ODP), including registration-free RDMA memory access
- > Enhanced congestion control
- > Burst buffer offload

In-Network Computing

- > Collective operations offloads
- > Vector collective operations offloads
- > MPI tag matching
- > MPI_Alltoall offloads
- > Rendezvous protocol offload
- > In-network memory

Hardware-Based IO Virtualization

- > Single root IO virtualization (SR-IOV)

Storage Offloads

- > Block-level encryption: XTS-AES 256/512-bit key
- > NVMe over Fabrics (NVMe-oF) offloads for target machine

- > T10 Data Integrity Field (DIF) signature handover operation at wire speed for ingress and egress traffic
- > Storage protocols: SRP, iSER, NFS RDMA, SMB Direct, NVMe-oF

HPC Software Libraries

- > NVIDIA HPC-X® and UCX®, UCC, NCCL, OpenMPI, MVAPICH, MPICH, OpenSHMEM, PGAS, and various commercial packages

Management and Control

- > NC-SI, MCTP over SMBus, and MCTP over PCIe
- > PLDM for Monitor and Control DSP0248
- > PLDM for Firmware Update DSP0267
- > PLDM for Redfish Device Enablement DSP0218
- > PLDM for FRU DSP0257
- > SPDMP DSP0274
- > General-purpose IO pins
- > Serial Peripheral Interface (SPI) to flash
- > JTAG 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)

Security

- > Secure boot with hardware root of trust
- > Secure firmware update
- > Flash encryption

Advanced Timing and Synchronization

- > Advanced PTP
- > IEEE 1588v2 (any profile)
- > Line-rate hardware timestamp (UTC format)
- > Configurable PPS In and configurable PPS Out
- > Time-triggered scheduling
- > PTP-based packet pacing
- > Time-Sensitive Networking (TSN)

Compatibility

PCI Express Interface

- > PCIe Gen 5.0 compatible, 32 lanes
- > Support for PCIe x1, x2, x4, x8, and x16 configurations
- > NVIDIA Multi-Host™ supports connection of up to 4 hosts
- > PCIe Atomic
- > Transaction layer packet (TLP) processing hints (TPH)
- > PCIe switch Downstream Port Containment (DPC)
- > Advanced error reporting (AER)
- > Access Control Service (ACS) for peer-to-peer secure communication
- > Process Address Space ID (PASID)
- > Address translation services (ATS)
- > Support for MSI/MSI-X mechanisms
- > Support for SR-IOV

Operating Systems/Distributions*

- > In-box drivers for major operating systems:
 - > Linux: RHEL, Ubuntu
 - > Windows
- > Virtualization and containers
 - > VMware ESXi (SR-IOV)
 - > Kubernetes
- > OpenFabrics Enterprise Distribution (OFED)
- > OpenFabrics Windows Distribution (WinOF-2)

*This section describes hardware features and capabilities.
Please refer to the driver and firmware release notes for feature availability.

Adapter Card Portfolio and Ordering Information

PCIe Standup Form Factor

InfiniBand Supported Speeds [Gb/s]	Network Ports and Cages	Host Interface [PCIe]	Form Factor	Orderable Part Number (OPN)
Up to 400Gb/s	1x OSFP	PCIe Gen 4.0/5.0 x16 With option for extension	HHHL	MCX75510AAS-NEAT
	1x OSFP	PCIe Gen 4.0/5.0 x16	HHHL	MCX75310AAS-NEAT
	1x OSFP	PCIe Gen 4.0/5.0 2x8 in a row	HHHL	MCX75210AAS-NEAT
Up to 200Gb/s	1x OSFP	PCIe Gen 4.0/5.0 x16 With option for extension	HHHL	MCX75510AAS-HEAT
	1x OSFP	PCIe Gen 4.0/5.0 x16	HHHL	MCX75310AAS-HEAT
	1x OSFP	PCIe Gen 4.0/5.0 2x8 in a row	HHHL	MCX75210AAS-HEAT
Up to 200Gb/s	2x QSFP	PCIe Gen 4.0/5.0 x16 With option for extension	HHHL	MCX755106AS-HEAT ¹

Dimensions without brackets are 167.65mm x 68.90mm. All adapters are shipped with the tall bracket mounted and a short bracket as an accessory.

Adapter card OPNs provided above support secure boot with hardware root of trust.

¹This card supports one port of InfiniBand, and a second port as either InfiniBand or Ethernet.

Auxiliary Boards

Passive Auxiliary Kit - Extension for 2x PCIe Gen 4.0 x16	Orderable Part Number (OPN)
15cm harness	MTMK9100-T15
35cm harness	MTMK9100-T35
55cm harness	MTMK9100-T55

OCP 3.0 Small Form Factor

InfiniBand Supported Speeds [Gb/s]	Network Ports and Cages	Host Interface [PCIe]	Form Factor	OPN
Up to 400Gb/s	1x OSFP	PCIe Gen 4.0/5.0 x16	TSFF	MCX75343AAS-NEAC ¹
Up to 200Gb/s	2x QSFP	PCIe Gen 4.0/5.0 x16	SFF	MCX753436AS-HEAB

The last digit of the OPN-suffix displays the default bracket option: B / C = pull tab, I / J = internal lock; E / F = ejector latch.

For other bracket types, contact NVIDIA.

Note 1. Pre OCP3.2 Spec

IC Ordering Information

Product Description	OPN
ConnectX-7 2-port IC, 400Gb/s, PCIe 5.0 x32, No Crypto	MT29108A0-NCCF-NV
ConnectX-7 2-port IC, 400Gb/s, Multi-Host, PCIe 5.0 x32, No Crypto	MT29108A0-NCCF-NVM
ConnectX-7 2-port IC, 400Gb/s, Multi-Host, PCIe 5.0 x32, Crypto	MT29108A0-CCCF-NVM

[Learn more](#)

Find out more at **NVIDIA ConnectX-7 InfiniBand**

© 2021 NVIDIA Corporation & Affiliates. All rights reserved. NVIDIA, the NVIDIA logo, ConnectX, GPUDirect, HPC-X, Mellanox, NVIDIA Multi-Host, NVIDIA Socket Direct, and UCX are trademarks and/or registered trademarks of NVIDIA Corporation and its affiliates in the U.S. and other countries. All other trademarks are property of their respective owners. DEC21/60456PB-R5.

