




NVIDIA QUANTUM-2 INFINIBAND PLATFORM

Extreme Performance for Exascale AI



The NVIDIA® Quantum-2 InfiniBand platform provides AI developers and scientific researchers with the highest networking performance available to take on the world's most challenging problems. New NVIDIA In-Network Computing acceleration engines provide ultra-low latency and double the data throughput, while delivering the scalability and feature-rich capabilities required for supercomputers, artificial intelligence, and hyperscale cloud data centers.

NVIDIA Quantum-2 enhances and extends In-Network Computing acceleration technology with preconfigured and programmable engines such as NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™, Message Passing Interface (MPI) tag matching, MPI_Alltoall, and programmable cores, as well as full transport offload with RDMA, GPUDirect RDMA and GPUDirect Storage, which deliver the best cost per node and return on investment (ROI).

World-Leading Performance

The NVIDIA Quantum-2 InfiniBand platform continues to set world records for high-performance networking, delivering 400Gb/s per port — 2X higher bandwidth compared to the previous generation, 3X higher switch silicon port density, 5X higher switch system capacity, 32X higher AI acceleration power per switch, and the ability to connect over one million nodes at 400Gb/s in a 3-hop Dragonfly+ topology.

Featuring the third generation of NVIDIA SHARP technology, the platform allows virtually unlimited scalability for large data aggregation for small and large data aggregations through the network — 32X higher AI acceleration power compared to the previous generation. Moreover, the third generation of SHARP technology enables multiple tenants or multiple parallel applications to share the infrastructure without any performance degradation. MPI_Alltoall acceleration and MPI tag matching hardware engines, and other features like advanced congestion control, adaptive routing, and self-healing networking provide critical enhancements to high-performance computing (HPC) and AI clusters, enabling them to reach even higher levels of performance than before.

NVIDIA Quantum-2 Portfolio

The NVIDIA Quantum-2 switch application-specific integrated circuit (ASIC) delivers 64 400Gb/s InfiniBand ports or 128 200Gb/s InfiniBand ports, the third generation of NVIDIA SHARP technology, and other InfiniBand features. The NVIDIA Quantum-2 platform includes fixed-configuration and modular switch systems based on the NVIDIA Quantum-2 switch ASIC.

The NVIDIA ConnectX[®]-7 InfiniBand host channel adapter (HCA) ASIC delivers 400Gb/s data throughput and supports 32 lanes of PCIe Gen5 or Gen4 for host connectivity. Built on the most advanced 100Gb/s-per-lane serializer/deserializer (SerDes) technology, 400Gb/s InfiniBand physical connectivity is based on octal small form-factor pluggable (OSFP) connectors on both the switches and HCA endpoints. Each switch OSFP connector holds two 400Gb/s InfiniBand ports or four 200Gb/s InfiniBand ports. Each HCA OSFP connector carries a single 400Gb/s InfiniBand port. The 400Gb/s cabling offering includes active and passive copper cables, transceivers, and multi-fiber push on (MPO) optical cables.

Fixed-Configuration Switches

The NVIDIA Quantum-2 family of fixed configuration switches consists of switches with 64 400Gb/s ports on 32 physical OSFP connectors, which can be split to deliver up to 128 200Gb/s ports. The compact 1U switches offering includes internally managed and externally managed (aka unmanaged) versions. The switches carry an aggregated bidirectional throughput of 51.2 terabits per second (Tb/s), with a landmark capacity of more than 66.5 billion packets per second. An ideal rack-mounted InfiniBand solution, the fixed configuration switch allows maximum flexibility as it enables a variety of topologies, including Fat Tree, DragonFly+, multi-dimensional Torus, and more.

Modular Switches

The NVIDIA Quantum-2 modular switch offering includes several flavors: 2,048 400Gb/s InfiniBand ports (that can be split into 4,096 200Gb/s InfiniBand ports); 1,024 400Gb/s ports (that can be split into 2,048 200Gb/s ports); or 512 400Gb/s ports (that can be split into 1024 200Gb/s ports). The large modular switch is based on a non-blocking, two-level Fat Tree “fabric-in-a-rack” occupying a three rack width. It carries a total unidirectional throughput of 819Tb/s or bidirectional throughput of 1.64 petabits per second (Pb/s), which is 5X over the previous generation of InfiniBand modular switches.

Host Channel Adapters

ConnectX-7 HCAs are offered in various form factors, delivering single or dual ports at 400Gb/s speeds with OSFP connectors or 200Gb/s speeds with quad small form-factor pluggable (QSFP112) connectors. The popular form factors meet the Card Electromechanical (CEM) Specification with 16 lanes of PCIe Gen5 or Gen4. Some versions come with the option to connect an additional 16-lane auxiliary card, which leverages NVIDIA Socket Direct[®] technology, to achieve 32 lanes of PCIe Gen4. Other form factors include Open Compute Project (OCP) 3.0 with OSFP connectors, OCP 3.0 with QSFP112 connectors, and CEM PCIe x16 with QSFP112 connectors.

The ConnectX-7 HCAs deliver advanced In-Network Computing with MPI_Alltoall and MPI tag matching hardware engines, and other fabric enhancement features such as quality of service (QoS), congestion control, and more.

Transceivers and Cables

400Gb/s InfiniBand (NDR) connectivity provides the maximum flexibility to build a topology of choice, using connectorized transceivers with passive fiber cables, active copper cables (ACCs), and direct attached copper cables (DACs). The usage of MPO fiber cables enables building complex infrastructures and for simpler installation, maintenance, and upgrades.

- > The connector plugs/cages for 400Gb/s InfiniBand are the 8-channel and 4-channel octal small form-factor plug (OSFP) and the 4-channel, quad small form-factor plug 112G (QSFP112).
- > Twin-port OSFP transceivers have two separate 400Gb/s NDR transceiver engines inside with two 4-channel, MPO/APC optical connectors. Both NDR connections can link to other switches as a pair at 800Gb/s aggregate, or as two separate 400Gb/s links to two switches, or two HCA/DPU adapter end points.
- > Each 4-channel NDR port can be split into two 2-channel, 200Gb/s NDR200 ports using fiber splitter cables; this enables linking one switch cage device to four HCA/DPU end points.
- > Similarly, switch-side, twin-port, copper DACs and ACCs are also offered in straight or 1:2 or 1:4 splitter cables with single-port OSFP or QSFP112 end points to HCA and DPUs.
- > The air-cooled Quantum-2 switches only accept twin-port, finned-top, OSFP devices. Liquid-cooled switches only accept twin-port, flat-top, OSFP devices. HCAs and DPUs only use flat-top OSFP or QSFP112 devices.

The large selection of device combinations creates the following array of interconnect choices:

- > Switch-side, twin-port DACs, ACCs, and optics connect to 400Gb/s NDR or 200Gb/s NDR200, single-port OSFP or QSFP112 devices in ConnectX-7 HCAs and BlueField-3 DPUs.
- > Connectivity to previous InfiniBand generations, switch-side, twin-port OSFP, 2xHDR DACs and AOCs splitter cables with QSFP56 connector end points enable two HDR, two EDR or four HDR100 links to Quantum switches, ConnectX-6 HCAs, or BlueField-2 DPUs.
- > Flat-top, twin-port transceivers and ACC cables can connect DGX-H100 “Hopper” GPU systems to Quantum-2 switches creating 400Gb/s NDR GPU compute fabrics.

UFM Cyber-AI

The NVIDIA Unified Fabric Manager (UFM) Cyber-AI platform offers enhanced and real-time network telemetry, together with AI-powered intelligence and advanced analytics. Enabling IT managers to discover operational anomalies and predict network failures, it improves security and data center uptime and decreases overall operating expenses.

Ordering Information

NVIDIA Quantum-2 Modular Switches

Orderable Part Number (OPN)	Description
MCS9500	NVIDIA 1,600b/s, 2,048-port 400Gb/s InfiniBand chassis and internal connectivity
MCS9510	NVIDIA 800Tb/s, 1,024-port 400Gb/s InfiniBand chassis and internal connectivity
MCS9520	NVIDIA 400Tb/s, 512-port 400Gb/s InfiniBand chassis and internal connectivity
MCS9505-AHX	NVIDIA MCS95xx modular systems liquid-to-air heat exchanger
MCS9500-KIT-AHX	NVIDIA MCS9500, 2 AHXs equipment kit
MCS9500-REDKIT-AHX	NVIDIA MCS9500, 3 AHXs redundancy equipment kit
MCS9510-KIT-AHX	NVIDIA MCS9510, 1 AHX equipment kit
MCS9510-REDKIT-AHX	NVIDIA MCS9510, 2 AHXs redundancy equipment kit
MQM9510-N	NVIDIA Quantum-2 400Gb/s InfiniBand, 2U leaf blade, 128 NDR ports, 64 OSFP ports
MQM9520-N	NVIDIA Quantum-2 400Gb/s InfiniBand, 2U spine blade, 128 NDR ports
MTDF-LIQ-D	NVIDIA MCS95xx chassis 19 liters PG25 coolant
MMB9500	NVIDIA MCS95xx management switch
MTDF-PDU-A	NVIDIA MCS95xx modular system PDU
GPS-CS9500-OST	CS9500 (2,048-port) 400Gb/s chassis installation, including onsite rack and liquid cooling system deployment. Price per chassis.
GPS-CS9510-20-OST	CS9510 (1,024-port) 400Gb/s chassis installation, including onsite rack and liquid cooling system deployment. Price per chassis.

NVIDIA Quantum-2 Fixed-Configuration Switches

Orderable Part Number (OPN)	Description
MQM9790-NS2F	NVIDIA Quantum-2-based 400Gb/s InfiniBand switch, 64 400Gb/s ports, 32 OSFP ports, non-blocking switching capacity of 51.2Tb/s, two power supplies (AC), standard depth, unmanaged, power-to-connector (P2C) airflow, rail kit
MQM9790-NS2R	NVIDIA Quantum-2-based 400Gb/s InfiniBand switch, 64 400Gb/s ports, 32 OSFP ports, non-blocking switching capacity of 51.2Tb/s, two power supplies (AC), standard depth, unmanaged, connector-to-power (C2P) airflow, rail kit
MQM9700-NS2F	NVIDIA Quantum-2-based 400Gb/s InfiniBand switch, 64 400Gb/s ports, 32 OSFP ports, non-blocking switching capacity of 51.2Tb/s, two power supplies (AC), standard depth, managed, P2C airflow, rail kit
MQM9700-NS2R	NVIDIA Quantum-2-based 400Gb/s InfiniBand switch, 64 400Gb/s ports, 32 OSFP ports, non-blocking switching capacity of 51.2Tb/s, two power supplies (AC), standard depth, managed, C2P airflow, rail kit

ConnectX-7 Host Channel Adapter (HCA)

PCIe Standup Adapters

OPN	Description
MCX75510AAS-NEAT	NVIDIA ConnectX-7 adapter card, 400Gb/s InfiniBand, single-port OSFP, Socket Direct ready, PCIe 5.0 x16 with IPEX connectors for extension, no crypto, tall bracket
MCX75310AAS-NEAT	NVIDIA ConnectX-7 adapter card, 400Gb/s InfiniBand, single-port OSFP, PCIe 5.0 x16, no crypto, tall bracket
MCX75510AAS-HEAT	NVIDIA ConnectX-7 adapter card, 200Gb/s, single-port OSFP, Socket Direct ready, PCIe 5.0 x16 with IPEX connectors for extension, no crypto, tall bracket
MCX75310AAS-HEAT	NVIDIA ConnectX-7 adapter card, 200Gb/s, single-port OSFP, PCIe 5.0 x16, no crypto, tall bracket
MCX75210AAS-NEAT	NVIDIA ConnectX-7 adapter card, 400Gb/s InfiniBand, single-port OSFP, PCIe 5.0 2x8 in a row, no crypto, tall bracket
MCX75210AAS-HEAT	NVIDIA ConnectX-7 adapter card, 200Gb/s, single-port OSFP, PCIe 5.0 2x8 in a row, no crypto, tall bracket
MCX755106AS-HEAT ¹	NVIDIA 200Gb/s dual-port virtual protocol interconnect (VPI), QSFP, PCIe Gen5 x16 HHHL, extension option

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

ConnectX-7 Host Channel Adapter (HCA) (continued)

Auxiliary Cards

OPN	Description
MTMK9100-T15	NVIDIA auxiliary kit for additional PCIe Gen4 x16 connection, PCIe Gen4 x16 passive auxiliary card, two 150 millimeter (mm) IPEX cables
MTMK9100-T25	NVIDIA auxiliary kit for additional PCIe Gen4 x16 connection, PCIe Gen4 x16 passive auxiliary card, two 250mm IPEX cables
MTMK9100-T35	NVIDIA auxiliary kit for additional PCIe Gen4 x16 connection, PCIe Gen4 x16 passive auxiliary card, two 350mm IPEX cables
MTMK9100-T55	NVIDIA auxiliary kit for additional PCIe Gen4 x16 connection, PCIe Gen4 x16 passive auxiliary card, two 550mm IPEX cables

Open Compute Project (OCP) Adapters

OPN	Description
MCX75343AAS-NEAC	NVIDIA 400Gb/s single port, OSFP, PCIe Gen5 x16 OCP3.0 ² small form factor (TSFF)
MCX753436AS-HEAB	NVIDIA 200Gb/s VPI dual port, QSFP, PCIe Gen5 x16 OCP3.0 (SFF)

² Pre OCP3.2 Specification

Transceivers and Cables

Transceivers for Air-Cooled Quantum-2 Switches

OPN	Description
Multimode Transceivers	
MMA4Z00-NS	NVIDIA twin-port 850nm transceiver, 800Gb/s, 2x NDR, 2x SR4, finned-top OSFP, 2x MPO/APC, up to 30m (switch-side, air-cooled)
MMA4Z00-NS-FLT	NVIDIA twin-port 850nm transceiver, 800Gb/s, 2x NDR, 2x SR4, flat-top OSFP, 2x MPO/APC, up to 30m (switch-side, liquid-cooled)
MMA4Z00-NS400	NVIDIA single-port 850nm transceiver, 400Gb/s, NDR, SR4, flat-top OSFP, MPO/APC, up to 30m (switch #2/HCA/DPU-side)
MMA4Z00-NS200	NVIDIA single-port 850nm transceiver, 200Gb/s, NDR200, SR2, flat-top OSFP, MPO/APC, up to 30m (HCA/DPU-side)
MMA1Z00-NS400	NVIDIA single-port 850nm transceiver, 400Gb/s, NDR, SR4, flat-top QSFP112, MPO/APC, up to 30m (HCA/DPU-side)
MMA1Z00-NS200	NVIDIA single-port 850nm transceiver, 200Gb/s, NDR200, SR2, flat-top QSFP112, MPO/APC, up to 30m (HCA/DPU-side)
Single-Mode Transceivers	
MMS4X00-NM	NVIDIA twin-port 1310nm transceiver, 800Gb/s, 2x NDR, 2x DR4, finned-top OSFP, 2x MPO/APC, up to 500m (switch-side, air-cooled)
MMS4X00-NS	NVIDIA twin-port 1310nm transceiver, 800Gb/s, 2x NDR, 2x DR4, finned-top OSFP, 2x MPO/APC, up to 100m (switch-side, air-cooled)
MMS4X00-NL	NVIDIA twin-port 1310nm transceiver, 800Gb/s, 2x NDR, 2x DR4, finned-top OSFP, 2x MPO/APC, up to 30m (switch-side, air-cooled)
MMS4X00-NM-FLT	NVIDIA twin-port 1310nm transceiver, 800Gb/s, 2x NDR, 2x DR4, flat-top OSFP, 2x MPO/APC, up to 500m (switch-side, liquid-cooled)
MMS4X00-NS-FLT	NVIDIA twin-port 1310nm transceiver, 800Gb/s, 2x NDR, 2x DR4, flat-top OSFP, 2x MPO/APC, up to 100m (switch-side, liquid-cooled)
MMS4X00-NL-FLT	NVIDIA twin-port 1310nm transceiver, 800Gb/s, 2x NDR, 2x DR4, flat-top OSFP, 2x MPO/APC, up to 30m (switch-side, liquid-cooled)
MMS4X00-NS400	NVIDIA single-port 1310nm transceiver, 400Gb/s, NDR, DR4, flat-top OSFP, MPO/APC, up to 100m (switch #2/HCA/DPU-side)
MMS4X00-NL400	NVIDIA single-port 1310nm transceiver, 400Gb/s, NDR, DR4, flat-top OSFP, MPO/APC, up to 30m (switch #2/HCA/DPU-side)
MMS4X00-NS200	NVIDIA single-port 1310nm transceiver, 200Gb/s, NDR200, DR2, flat-top OSFP, MPO/APC, up to 100m (HCA/DPU-side)
MMS4X00-NL200	NVIDIA single-port 1310nm transceiver, 200Gb/s, NDR200, DR2, flat-top OSFP, MPO/APC, up to 30m (HCA/DPU-side)

Transceivers and Cables (continued)

MPO/APC Single-Mode Crossover Fiber

MFP7E30-Nxxx	NVIDIA passive fiber cable, SMF, MPO to MPO xxx indicates length in meters: 001, 002, 003, 005, 007, 010, 015, 020, 030, 040, 050, 060, 070, 100, 150
--------------	--

MPO/APC Single-Mode Crossover Fiber Splitter

MFP7E40-Nxxx	NVIDIA passive fiber cable, SMF, MPO to 2x MPO xxx indicates length in meters: 003, 005, 007, 010, 015, 020, 030, 040, 050
--------------	---

MPO/APC Multimode Crossover Fiber

MFP7E10-Nxxx	NVIDIA passive fiber cable, MMF, MPO to MPO xxx indicates length in meters: 003, 005, 007, 010, 015, 020, 030
--------------	--

MPO/APC Multimode Crossover Fiber Splitter

MFP7E20-Nxxx	NVIDIA passive fiber cable, MMF, MPO to 2x MPO xxx indicates length in meters: 003, 005, 007, 010, 015, 020, 030
--------------	---

Direct Attached Copper (DAC) Switch to Switch OSFP

MCP4Y10-Nxxx	NVIDIA passive copper cable, InfiniBand NDR 800Gb/s to 800Gb/s, OSFP to OSFP xxx indicates length in meters: 00A (0.5m), 00B (0.75m), 001, 002
--------------	---

Direct Attached Copper (DAC) Switch to ConnectX-7 HCA OSFP

MCP7Y00-Nxxx	NVIDIA passive copper splitter cable, InfiniBand NDR 800Gb/s to 2x 400Gb/s, OSFP to 2x OSFP xxx indicates length in meters: 001, 01A (1.5m), 002, 02A (2.5m), 003
MCP7Y50-Nxxx	NVIDIA passive copper splitter cable, InfiniBand NDR 800Gb/s to 4x 200Gb/s, OSFP to 4x OSFP xxx indicates length in meters: 001, 01A (1.5m), 002, 02A (2.5m), 003

Direct Attached Copper (DAC) Switch to BlueField-3 DPU QSFP112

MCP7Y10-Nxxx	NVIDIA passive copper splitter cable, InfiniBand NDR 800Gb/s to 2x 400Gb/s, OSFP to 2x QSFP112 xxx indicates length in meters: 001, 01A (1.5m), 002, 02A (2.5m), 003
--------------	---

Direct Attached Copper (DAC) Switch to ConnectX-7 HCA or BlueField-3 DPU QSFP112

MCP7Y40-Nxxx	NVIDIA passive copper splitter cable, InfiniBand NDR 800Gb/s to 4x 200Gb/s, OSFP to 4x QSFP112 xxx indicates length in meters: 001, 01A (1.5m), 002, 02A (2.5m), 003
--------------	---

Active Copper (ACC) Switch to Switch OSFP

MCA4J80-Nxxx	NVIDIA active copper cable, InfiniBand NDR 800Gb/s to 800Gb/s, OSFP to OSFP xxx indicates length in meters: 003, 004, 005
--------------	--

Active Copper Cable (ACC) Switch to ConnectX-7 HCA OSFP

MCA7J60-Nxxx	NVIDIA active copper splitter cable, InfiniBand NDR 800Gb/s to 2x 400Gb/s, OSFP to 2x OSFP xxx indicates length in meters: 004, 005
MCA7J70-Nxxx	NVIDIA active copper splitter cable, InfiniBand NDR 800Gb/s to 4x 200Gb/s, OSFP to 4x OSFP xxx indicates length in meters: 004, 005

Transceivers and Cables (continued)

Active Copper Cable (ACC) Switch to BlueField-3 DPU QSFP112

MCA7J65-Nxxx	NVIDIA active copper splitter cable, InfiniBand NDR 800Gb/s to 2x 400Gb/s, OSFP to 2x QSFP112 xxx indicates length in meters: 004, 005
MCA7J75-Nxxx	NVIDIA active copper splitter cable, InfiniBand NDR 800Gb/s to 4x 200Gb/s, OSFP to 4x QSFP112 xxx indicates length in meters: 004, 005

Active Copper Cable (ACC) Switch to ConnectX-7 HCA or BlueField-3 DPU QSFP112

MCA7J75-Nxxx	NVIDIA active copper splitter cable, InfiniBand NDR 800Gb/s to 4x 200Gb/s, OSFP to 4x QSFP112 xxx indicates length in meters: 004, 005
--------------	---

Backward-Compatible Active Optical Cable (AOC), 1:2 Splitter Switch to ConnectX-6, BlueField-2, or Quantum Switch - HDR

MFA7U10-Hxxx	NVIDIA AOC 1:2 splitter, InfiniBand 2x HDR to 2x HDR 400Gb/s to 2x 200Gb/s, OSFP to 2x QSFP56 xxx indicates length in meters: 003, 005, 010, 015, 020, 030, 050
--------------	--

Backward-Compatible Active Optical Cable (AOC), 1:4 Splitter Switch to ConnectX-6, BlueField-2, or Quantum Switch - HDR100

MFA7U20-Hxxx	NVIDIA AOC 1:4 splitter, InfiniBand 2x HDR to 4x HDR100 400Gb/s to 4x 100Gb/s, OSFP to 4x QSFP56 xxx indicates length in meters: 003, 005, 010, 015, 020, 030, 050
--------------	---

Backward-Compatible Active Optical Cable (AOC), 1:2 Splitter Switch to ConnectX-5, BlueField-2, Quantum, or Switch-IB 2 - EDR

MFA7U30-Exxx	NVIDIA AOC 1:2 splitter, InfiniBand 2x EDR to 2x EDR 200Gb/s to 2x 100Gb/s, OSFP to 2x QSFP28 xxx indicates length in meters: 003, 005, 010, 015, 020, 030, 050
--------------	--

Backward-Compatible Direct Attach Copper Cable (DAC), 1:2 Splitter Switch to ConnectX-6, BlueField-2, or Quantum Switch - HDR

MCP7Y60-Hxxx	NVIDIA DAC 1:2 splitter, InfiniBand 2x HDR to 2x HDR 400Gb/s to 2x 200Gb/s, OSFP to 2x QSFP56 xxx indicates length in meters: 001, 01A (1.5m), 002
--------------	---

Backward-Compatible Direct Attach Copper Cable (DAC), 1:4 Splitter Switch to ConnectX-6, BlueField-2, or Quantum Switch - HDR100

MCP7Y70-Hxxx	NVIDIA DAC 1:4 splitter, InfiniBand 2x HDR to 4x HDR100, 400Gb/s to 4x 100Gb/s, OSFP to 4x QSFP56 xxx indicates length in meters: 001, 01A (1.5m), 002
--------------	---

Backward-Compatible Direct Attach Copper Cable (DAC), 1:2 Splitter Switch to ConnectX-6, BlueField-2, or Quantum Switch - EDR

MCP7Y60-Hxxx	NVIDIA DAC 1:2 splitter, InfiniBand 2x EDR to 2x EDR 200Gb/s to 2x 100Gb/s, OSFP to 2x QSFP28 xxx indicates length in meters: 001, 01A (1.5m), 002
--------------	---

[Learn more](#)

To learn more about NVIDIA Quantum-2 InfiniBand platform, visit [NVIDIA.com/InfiniBand](https://www.nvidia.com/InfiniBand).

© 2022 NVIDIA Corporation & Affiliates. All rights reserved. NVIDIA, the NVIDIA logo, ConnectX, Scalable Hierarchical Aggregation and Reduction Protocol (SHARP), and Socket Direct 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. MAY22

