



NVIDIA QUANTUM PRODUCT GUIDE

End-to-End High Performance InfiniBand Networking for HPC, AI, and Cloud Data Centers



Boost the Availability and Scale of High Performance Computing (HPC) and AI Data Centers for Deep Learning, Research, Innovation, and Product Development

Complex workloads demand ultra-fast processing of high resolution simulations, extreme-size data sets, and highly parallelized algorithms. As these computing requirements continue to grow, NVIDIA® Quantum InfiniBand platform—the world's only fully offloadable, In-Network Computing interconnect technology—provides the dramatic leap in performance needed to achieve unmatched performance in HPC, AI, and hyperscale cloud infrastructures, with less cost and complexity. NVIDIA Quantum-2, the seventh generation of NVIDIA InfiniBand architecture, gives AI developers, scientific researchers, and product designers even faster networking performance and feature sets to take on the world's most challenging problems.

NVIDIA Quantum Switches and Routers

NVIDIA Quantum InfiniBand switch systems deliver the highest performance and port density available. Innovative In-Network Computing capabilities, such as NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ and advanced management features, such as self-healing network capabilities, quality of service, enhanced congestion control provide a performance boost for industrial, AI, and scientific applications.


In addition, NVIDIA InfiniBand routers enable subnet isolation, while maintaining the high performance that is critical to large-scale and diverse data centers. The extensive switch portfolio enables compute clusters to operate at any scale, while reducing capital expenses, operational costs, and infrastructure complexity.

NVIDIA BlueField Data Processing Units (DPUs)

The NVIDIA BlueField® DPUs combine powerful computing, high-speed networking, and extensive programmability to deliver software-defined, hardware-accelerated solutions for the most demanding workloads. By offloading, accelerating, and isolating a broad range of advanced networking, storage, and security services, BlueField DPUs provide a secure and accelerated infrastructure. From accelerated AI computing to cloud-native supercomputing, BlueField redefines what's possible.

NVIDIA ConnectX Smart Host Channel Adapters (HCAs)

NVIDIA ConnectX® InfiniBand adapters provide ultra-low latency, extreme throughput, and innovative In-Network Computing engines such as MPI Tag Matching and Alltoall hardware engines, to deliver the acceleration, scalability, and feature-rich technology needed for today's and tomorrow's modern workloads.



NVIDIA MetroX Long Haul Switches

NVIDIA MetroX® systems extend the reach of InfiniBand to up to 40 kilometers, enabling native RDMA connectivity between remote data centers, or between data centers and remote storage, for high availability and disaster recovery. Delivering up to 100Gb/s data throughput on the long-haul ports, and 200Gb/s on standard ports, MetroX-2 users can easily migrate application jobs from one InfiniBand center to another, or to combine the compute power of multiple remote data centers together for higher overall performance and scalability.




NVIDIA Skyway Gateway to Ethernet

The NVIDIA Skyway™ InfiniBand to Ethernet gateway appliance enables scalable and efficient connectivity from high performance, low-latency InfiniBand data centers to external Ethernet networks and infrastructures. Supporting high availability and load balancing with standard IP routing protocols, Skyway empowers InfiniBand-based high performance and cloud data centers to achieve the lowest interconnect latency, while also providing a simple and cost-effective option to connect to external Ethernet networks.



NVIDIA LinkX Cables and Transceivers

NVIDIA LinkX® InfiniBand cables and transceivers maximize the performance of InfiniBand networks to deliver high-bandwidth, low-latency, highly reliable, and robust connectivity. To provide superior system performance, NVIDIA ensures the highest quality in all LinkX products.



NVIDIA UFM Fabric and Network Management

NVIDIA UFM® (Unified Fabric Management) platforms revolutionize data center networking management. Supporting scale-out InfiniBand data centers, the UFM platforms combine enhanced, real-time network telemetry with AI-powered cyber intelligence and analytics, to realize higher utilization of fabric resources and a competitive advantage, while reducing OPEX. The UFM software suite includes, but is not limited to, fabric diagnostics, monitoring, alerting, provisioning, and advanced features such as congestion monitoring, and fabric segmentation and isolation.

UFM CYBER AI

UFM ENTERPRISE

UFM TELEMETRY



NVIDIA Acceleration Software

NVIDIA HPC-X® is a comprehensive software package that includes Message Passing Interface (MPI), Symmetrical Hierarchical Memory (SHMEM), Partitioned Global Address Space (PGAS) communications libraries, and various acceleration packages. This full-featured, tested, and packaged toolkit enables MPI and SHMEM/PGAS programming languages to achieve high performance, scalability, efficiency, and ensures that communication libraries are fully optimized by NVIDIA Quantum InfiniBand networking solutions.

NVIDIA INFINIBAND SWITCH SYSTEMS

| Switch | Configuration | Advanced Features | Size |
|--|---|--|--|
| QM9700 NVIDIA Quantum-2 Fixed Configuration Switch Series | 64 400Gb/s ports that can be split into 128 200Gb/s ports 51.2Tb/s total throughput Ultra-low switch latency | <ul style="list-style-type: none"> > Accelerated In-Network Computing > Limitless scalability > 3rd generation of SHARP (SHARPV3) > Enhanced management with out-of-the box bring-up for up to 2,000 nodes > Internally managed and externally-managed flavors | 1U |
| CS9500 NVIDIA Quantum-2 Modular Switch Series | Three systems included: 2,048 400Gb/s ports or 4,096 200Gb/s ports for 1,600Tb/s total throughput 1,024 400Gb/s ports or 2,048 NDR 200Gb/s ports for 800Tb/s switching capacity 512 400Gb/s port or 1,024 200Gb/s ports for 400Tb/s total throughput Ultra-low switch latency | <ul style="list-style-type: none"> > AI acceleration > SHARPV3: In-network collective offloads support low-latency and streaming aggregation for AI applications > Advanced adaptive routing, congestion control, and QoS > Self-healing networking > Liquid-cooled system | 48U Width: 2,048 ports - 3.5 std. racks (2100mm) 1,024 port - 2 std. racks (1200mm) 512 ports - 2 std. racks (1200mm) |
| QM8700 NVIDIA Quantum Fixed Configuration Switch Series | 40 200Gb/s ports or 80 100Gb/s ports 16Tb/s aggregate switch throughput Ultra-low switch latency | <ul style="list-style-type: none"> > Internally managed and externally managed flavors > Self-healing networking > SHARPV2: In-network collective offloads support of low-latency and streaming aggregation for AI applications > Adaptive routing, congestion control, and QoS | 1U |
| CS8500 NVIDIA Quantum Modular Switch Series | 800 200Gb/s ports or 1600 100Gb/s ports for 320Tb/s switching capacity Ultra-low switch latency | <ul style="list-style-type: none"> > Fabric with self-healing autonomy > SHARPV2: In-network collective offloads support of low-latency and streaming aggregation for AI applications, adaptive routing, congestion control, and QoS liquid-cooled system | 29U |
| SB7800 NVIDIA Switch-IB 2 Fixed Configuration Switch Series | 36 100Gb/s ports 7.2Tb/s aggregate switch throughput Ultra-low switch latency | <ul style="list-style-type: none"> > Internally managed and unmanaged flavors > Adaptive routing and QoS > SHARP: In-network collective offloads for low-latency HPC applications | 1U |
| CS7500 NVIDIA Switch-IB 2 Modular Switch Series | Three systems included: 648 100Gb/s ports for 130Tb/s switching capacity 324 100Gb/s for 65Tb/s switching capacity 216 100Gb/s for 43.2Tb/s switching capacity Ultra-low switch latency | <ul style="list-style-type: none"> > Adaptive routing and QoS > SHARP: In-network collective offloads for low-latency HPC applications > N+N power supply | 648 ports-28U 324 ports-16U 216 ports-12U |
| NVIDIA Skyway InfiniBand to Ethernet Gateway | 8 200Gb/s/100Gb/s InfiniBand ports 8 200/100Gb/s Ethernet ports 1.6Tb/s aggregate switch throughput | <ul style="list-style-type: none"> > Industry-leading InfiniBand to Ethernet gateway > Future-ready architecture | 1U |
| NVIDIA MetroX-2 Switch Systems | 2 100Gb/s InfiniBand QSFP28 long-haul ports 8 200Gb/s InfiniBand QSFP56 local ports | <ul style="list-style-type: none"> > Adaptive routing and congestion control > SHARP: In-network collective offloads for low-latency HPC applications > Self-healing networking | 1U |
| SB7880 NVIDIA Switch-IB 2 Router | 36 100Gb/s ports 7.2Tb/s aggregate switch throughput Up to 7.02 billion messages-per-second | <ul style="list-style-type: none"> > InfiniBand routing among different subnets > SHARP: In-Network collective offloads for low-latency HPC applications > Self-healing networking | 1U |

NVIDIA DATA PROCESSING UNITS (DPUS)

| DPU | Speeds | Connectors | Bus | Capabilities | Form Factor |
|--------------------|---------------|------------|-------------------|--|---------------|
| BlueField-3 | Up to 400Gb/s | QSFP112 | PCIe Gen3/4/5 x16 | BlueField-3 SoC <ul style="list-style-type: none"> > 16X Arm A78 cores > 400Gb/s InfiniBand ConnectX-7 hardware offloads | PCIe stand-up |
| BlueField-2 | Up to 200Gb/s | QSFP56 | PCIe Gen3/4 x16 | BlueField-2 SoC <ul style="list-style-type: none"> > 168x A72 Arm cores > 200Gb/s InfiniBand ConnectX-6 Dx hardware offloads | PCIe stand-up |

NVIDIA ADAPTERS

| Adapters (HCAs) | Speeds | Connectors | Bus | RDMA Message Rate (Mmps) | AdvancedFeatures | Form Factors |
|-----------------|---------------|-----------------|-------------------------------------|--|---|---|
| ConnectX-7 | Up to 400Gb/s | QSFP or QSFP112 | PCIe Gen5 x16 2x PCIe Gen4 x16 | 330-370 | <ul style="list-style-type: none"> > MPI_Alltoall offloads > Enhanced Congestion Control > Secure boot with hardware root of trust > NVIDIA Multi-Host up to x4 hosts (specific OPNs) > Advanced PTP > SHARP v3 support | PCIe stand-up PCIe Socket Direct OCP 3.0 |
| ConnectX-6 | Up to 200Gb/s | QSFP56 | PCIe Gen3/4 x16 2x PCIe Gen3 x16 | 215 | <ul style="list-style-type: none"> > Congestion Control > MPI tag matching offload > Block-level XTS-AES hardware encryption > Secure firmware update > NVIDIA Multi-Host up to 4 hosts > NVMe-oF Target Offload | PCIe stand-up PCIe Socket Direct OCP 3.0 |
| ConnectX-5 | Up to 100Gb/s | QSFP28 | PCIe Gen 3/4 x16 | 200 ConnectX-5 Ex (Gen4 server) 165 (Gen3 server) | <ul style="list-style-type: none"> > MPI tag matching offload > Host management > NVIDIA Multi-Host up to 4 hosts > T-10 DIF/Signature Handover | PCIe stand-up PCIe Socket Direct OCP 2.0 OCP 3.0 |

INFINIBAND INTERCONNECT

| Direct Attach Copper (DAC) | Active Copper Cables (ACC) | Active Optical Cables (AOC) | Optical Transceivers |
|--|--|---|--|
| <ul style="list-style-type: none"> > Maximum reaches: <ul style="list-style-type: none"> 400Gb/s: 1.5m 200Gb/s: 0.5m-2m 100Gb/s: 0.5m-5m > Zero power consumption > Near zero latency delays | <ul style="list-style-type: none"> > Maximum reaches: <ul style="list-style-type: none"> 400Gb/s: 1, 2, 3m includes 1:2 and 1:4 splitters for 400Gb/s and 200Gb/s 200Gb/s: 2, 3m includes 1:2 splitters > Booster IC extends copper reach > Maintains low-cost copper interconnects > Offers thinner cables than DACs | <ul style="list-style-type: none"> > Maximum reaches: <ul style="list-style-type: none"> 200Gb/s : 3m-150m 100Gb/s: 3m-30m 1:2 splitters 100Gb/s: 3m-100m > Lowest cost optical link > No optical connector cleaning > Increase reliability over transceivers | <ul style="list-style-type: none"> > Maximum reaches: <ul style="list-style-type: none"> 400Gb/s: Single mode Twin port 150m single mode and multi-mode 30m 200Gb/s: Single mode 2km; multi-mode 100m 100Gb/s: Single mode 40km; multi-mode 100m > Maximum reach optical links > Enables many configuration options |

Ordering Information

For ordering information, please contact gopny@pny.com

[Learn more](#)

Learn more at: pny.com/networking