



NVIDIA MELLANOX INFINIBAND PRODUCT GUIDE

NVIDIA Mellanox smart and high-speed InfiniBand interconnect technology solutions accelerate the world's leading supercomputing, artificial intelligence, and cloud platforms at HDR/EDR/FDR speeds. The end-to-end InfiniBand solution comprises a complete portfolio of network adapters, switches, cables and transceivers, gateways, fabric management and acceleration software suite, for compute and data-intensive applications. Backward and forward compatible, the InfiniBand solutions optimize data center efficiency and provide the best return on investment.

INFINIBAND SWITCHES AND ROUTERS


From NVIDIA Mellanox Switch-IB®-2 EDR 100Gb/s to NVIDIA Mellanox Quantum™ HDR 200Gb/s InfiniBand, the InfiniBand family of 1U and modular switches delivers the highest density, performance and scalability. The switches include the advanced In-Network Computing SHARP (Mellanox Scalable Hierarchical Aggregation and Reduction Protocol) technology. This technology enables performing data reductions on transferred network data, to accelerate high performance and deep learning applications.

InfiniBand switches also leverage the NVIDIA Self-Healing Networking technology to increase data center network resiliency by 5000 times compared to other software-based solution options. In addition, InfiniBand Routers enable new levels of subnet isolation and compute to storage connectivity, while maintaining high performance 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.



INFINIBAND ADAPTERS

Leveraging faster speeds and the innovative In-Network Computing, NVIDIA Mellanox HDR (200Gb/s), EDR (100Gb/s) and FDR (56Gb/s) InfiniBand adapters deliver the highest throughput and message rate in the industry. Providing best-in-class network performance, scale, and efficiency, these InfiniBand adapters enable extremely low-latency, and advanced application acceleration engines for high-performance computing, machine learning, cloud, storage, databases, and embedded applications, reducing cost-per-operation and increasing ROI.



INFINIBAND LONG HAUL

NVIDIA Mellanox MetroX®-2 systems extend the reach of InfiniBand to up to 40 kilometers, enabling native InfiniBand connectivity between remote data centers, or between data center and remote storage infrastructures, for high availability and disaster recovery. Delivering up to 100Gb/s data throughout long-haul systems, and 200Gb/s on standard links, MetroX-2 enables native RDMA connectivity, advanced routing, and scalability across distributed compute or storage platforms. MetroX-2 enables users to 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.




INFINIBAND GATEWAY TO ETHERNET

The NVIDIA Mellanox 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 future-ready architecture, Mellanox 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.




INFINIBAND CABLES AND TRANSCEIVERS

NVIDIA Mellanox LinkX® InfiniBand cables and transceivers maximize the performance of the 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.



INFINIBAND FABRIC AND NETWORK MANAGEMENT

NVIDIA Mellanox 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. Users can manage small to extremely large fabrics as a set of inter-related business entities while also performing fabric monitoring and optimizing performance at the application-logical level rather than only at the individual port or device level.



INFINIBAND ACCELERATION SOFTWARE

The NVIDIA Mellanox HPC-X® ScalableHPC Toolkit is a comprehensive MPI and SHMEM/PGAS software suite for high performance computing environments. For scientific research and engineering simulations, the complete HPC-X software toolkit, including SHARP, UCX, and other MPI acceleration engines, provides enhancements that significantly increase the scalability and performance of message communications in the network. HPC-X enables the rapid deployment and delivery of maximum application performance without the complexity and costs of licensed third-party tools and libraries.

INFINIBAND SYSTEMS

Switch	Performance	Advanced Features	Size
QM8700 Mellanox Quantum Switch Series	40X HDR (200Gb/s) ports 80X HDR100 (100Gb/s) ports 16Tb/s aggregate switch throughput	<ul style="list-style-type: none">> Internally managed and externally managed flavors> Self-Healing Networking> SHARP v2—In-network collective offloads support low-latency and streaming Aggregation for AI applications> Adaptive Routing, congestion control and QoS	1U
CS8500 Mellanox Quantum Director Series	Up to 800 HDR (200Gb/s) ports Up to 1600 HDR100 (100Gb/s) ports Up to 320Tb/s switching capacity Ultra-low switch latency	<ul style="list-style-type: none">> Fabric with self-healing autonomy> SHARP v2—In-network collective offloads support low-latency and streaming Aggregation for AI applications,> Adaptive routing, congestion control and QoS> Liquid-cooled system	29U

† For illustration only. Actual products may vary.

INFINIBAND SYSTEMS (CONT...)

Switch	Performance	Advanced Features	Size
SB7800 Mellanox Switch-IB® 2 Switch Series	36X EDR (100Gb/s) ports 7.2Tb/s aggregate switch throughput Up to 7.02 billion messages-per-second 90ns latency; 136W typical power consumption	<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 Mellanox Switch-IB® 2 Director Series	Up to 648 EDR (100Gb/s) Up to 130Tb/s switching capacity Ultra-low 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
Mellanox Skyway™ InfiniBand to Ethernet Gateway	8X HDR/HDR100/EDR ports 8X 200/100Gb/s Ethernet ports 1.6 Tb/s aggregate switch throughput	<ul style="list-style-type: none"> > Industry-leading InfiniBand to Ethernet gateway > Future-ready architecture 	2U
MetroX®-2 Switch Systems	2X EDR QSFP28 long-haul ports 8X HDR 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

INFINIBAND ADAPTERS AND SMART ADAPTERS

Adapters	Speeds	Connectors	Bus	RDMA Message Rate (mmpps)	Features	Form Factors
ConnectX-6	HDR HDR100 EDR FDR	QSFP56	PCIe Gen3/4 x16 2x PCIe Gen3 x16	215	<ul style="list-style-type: none"> > 0.6usec latency > Enhanced Congestion Control > MPI tag matching offload > Block-level XTS-AES hardware encryption Hairpin > Host management > NVIDIA Mellanox Multi-Host® > NVMe-oF Target Offload > T-10 Dif/Signature Handover 	PCIe stand-up PCIe Socke Direct OCP 3.0
ConnectX-5	EDR FDR	QSFP28	PCIe Gen 3/4 x16	200 (ConnectX-5 Ex Gen4 server) 165 (Gen3 server)	<ul style="list-style-type: none"> > 0.6usec latency > MPI tag matching offload > Host management > Mellanox Multi-Host > T-10 DIF/Signature Handover 	PCIe stand-up PCIe Socket Direct OCP 2.0
BlueField-2	HDR EDR	QSFP56	PCIe Gen3/4 x16	Contact Mellanox	<ul style="list-style-type: none"> > Secure hardware, isolation and cryptography > ConnectX-6 hardware offloads > NVIDIA Mellanox NVMe SNAP™ > 1GbE out-of-band management port for the Arm subsystem 	PCIe stand-up

CABLES AND TRANSCEIVERS

Direct Attach Cables (DAC)	Active Optical Cables (AOC)	Optical Transceivers
FDR/EDR/HDR copper cables Reach:0.5–2m (HDR), 0.5–5m (EDR) Zero power consumption Near zero latency delays	FDR/EDR/HDR optical link Reach: up to 100m Lowest power consumption: 2.2W (EDR) Low latency delays Low Smoke, Zero Halogen jacket (EDR and HDR) BER less than 1E-15	Full line of FDR/EDR/HDR 100m–2Km (HDR) 100m–40Km (EDR) SR4 lowest power consumption—2.2W (EDR) Multi-mode + single-mode

Learn more at www.mellanox.com/products/InfiniBand

© 2020 Mellanox Technologies. All rights reserved. NVIDIA, the NVIDIA logo, Mellanox, SHIELD, Scalable Hierarchical Aggregation and Reduction Protocol (SHARP), HPC-X, UCX, Mellanox Quantum, LinkX, BlueField, and ConnectX are trademarks and/or registered trademarks of Mellanox Technologies Ltd. and/or NVIDIA Corporation in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated. 60342BR R6/Nov20

