BROCHURE

InfiniBand Product Guide

End-to-end InfiniBand Solutions at HDR / EDR / FDR Speeds

Mellanox Technologies is the leading provider of the standard InfiniBand, smart high-speed interconnect technology solutions, accelerating the world's leading supercomputing, artificial intelligence and cloud platforms. Mellanox offers a complete InfiniBand portfolio of network adapters, switches, cables & transceivers, gateways, fabric management and Mellanox ScalableHPC® software suite, for compute and data intensive applications. Mellanox solutions are backward and forward compatible, optimizing data center efficiency and providing the best return on investment.

InfiniBand Switches

From Switch-IB®-2 100Gb/s EDR to Mellanox Quantum™ 200Gb/s HDR InfiniBand, the Mellanox family of 1RU and modular switches deliver the highest density and performance. Featuring the in-network computing Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ technology, Mellanox's extensive switch portfolio enables compute clusters to operate at any scale, while reducing operational costs and infrastructure complexity. InfiniBand switches also leverage Mellanox's SHIELD™ (Self-Healing Interconnect Enhancement for Intelligent Datacenters), which increases data center network resiliency by 5000 times, compared to other software-based solution options.

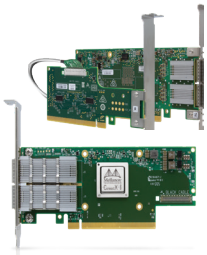

InfiniBand Adapters

Leveraging faster speeds and innovative In-Network Computing, Mellanox's award-winning 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, 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. Mellanox smart adapters deliver the highest performance and scalable connectivity for Intel, AMD, IBM Power, NVIDIA, Arm and FPGA-based Compute and Storage Platforms.

InfiniBand LongHaul

Mellanox's MetroX®-2 systems extend the reach of InfiniBand to up to 40 kilometers, enabling native InfiniBand connectivity between remote data centers, remote data center and storage infrastructures, or for disaster recovery. Delivering up to 100Gb/s data throughout, MetroX-2 enables native RDMA connectivity, high data throughput, advanced routing, and more, 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

Mellanox Skyway™ 200 Gigabit HDR InfiniBand to Ethernet gateway appliance enables scalable and efficient connectivity from high performance, low-latency InfiniBand data centers to external Ethernet networks and infrastructures. Mellanox Skyway™ empowers InfiniBand-based high performance and cloud data centers to achieve the lowest interconnect latency, while providing a simple and cost-effective option to connect to remote Ethernet networks.

InfiniBand Cables and Transceivers

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

InfiniBand Telemetry and Software Management

Mellanox's comprehensive suite of network telemetry and management software provides an innovative application-centric approach to bridge the gap between servers, applications and fabric elements. Mellanox's UFM® (Unified Fabric Management) software includes 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 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 Mellanox HPC-X software toolkit, including Mellanox UCX and FCA 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 Quantum Switch Series	40 X HDR (200Gb/s) ports 80 X HDR100 (100Gb/s) ports 16Tb/s aggregate switch throughput Up to 15.8 billion messages-per-second	Managed and unmanaged flavors SHIELD™ – Fabric with self-healing autonomy SHARP™ v2 – Collective offloads supporting streaming for Machine Learning Adaptive Routing, congestion control and QoS	1U
CS8500 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	SHIELD™ – Fabric with self-healing autonomy SHARP™ v2 – Collective offloads supporting streaming for Machine Learning Adaptive Routing; congestion control and QoS; Liquid-cooled system	29U
SB7800 Switch-IB® 2 Switch Series	36 X EDR (100Gb/s) ports 7.2Tb/s aggregate switch throughput Up to 7.02 billion messages-per-second 90ns latency; 136W typical power consumption	Managed and unmanaged flavors Adaptive routing and QoS SHARP™ – In-network collective offloads	1U
CS7500 Switch-IB® 2 Director Series	Up-to 648 EDR (100Gb/s) Up-to 130Tb/s switching capacity Ultra-low latency	648 ports, 324 ports and 216 ports flavors Adaptive routing and QoS SHARP™ – In-network collective offloads N+N power supply	648 ports - 28U 324 ports - 16U 216 ports - 12U
Mellanod Skyway™ InfiniBand to Ethernet Gateway	8 X HDR/HDR100/EDR ports 8 X 200/100Gb/s Ethernet ports 1.6 TB/s aggregate switch throughput	Industry leading InfiniBand to Ethernet gateway In-Network Computing technology delivers 10x higher performnance Future-ready to support 400Gb/s NDR and 800Gb/s XDR speeds	1U
MetroX®-2 LongHaul Switch Systems	2 X EDR QSFP28 long-haul ports 8 X HDR QSFP28 local ports 200Gb/s throughput	Adaptive Routing & Congestion Control SHARP™ – In-network collective offloads SHIELD™ – Fabric with self-healing autonomy	1U
Third-Party Tools and Libraries: HPC-X provides a compressive and high performance MPI and SHMEM communication libraries, including the support of Mellanox In-Network Computing hardware acceleration engines			

InfiniBand Smart Adapters & SmartNICs

Adapters	Speeds	Connectors	Bus	DMA Message Rate (mmps)	Features	Form Factors
ConnectX®-6	HDR HDR100 EDR FDR	QSFP56	PCIe Gen3/4 x16 2x PCIe Gen3 x16	215	0.6usec latency Enhanced Congestion Control MPI tag matching offload Block-level XTS-AES hardware encryption Hairpin (Host chaining) Host management, Mellanox Multi-Host® NVMe-of Target Offload Erasure Coding (RAID 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)	0.6usec latency MPI tag matching offload Host management Mellanox Multi-Host Erasure Coding (RAID Offload) 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	Secure hardware, isolation, cryptography ConnectX-6 hardware offloads NVMe SNAPTM 1GbE out-of-band management port for the Arm subsystem	PCIe stand-up

Cables & Transceivers

Direct Attach Cables (DAC)	Active Optical Cables (AOC)	Optical Transceivers
Cost-effective FDR/EDR/HDR cables Reach:0.5½ - 2m (HDR), 0.5½ - 5m (EDR) Zero power consumption Near zero latency delays	Cost-effective 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 IE-15	Full line of FDR/EDR/HDR 100m - 2Km (HDR) 100m - 40Km (EDR) SR4 lowest power consumption - 2.2W (EDR) Multi-mode + single-mode

350 Oakmead Parkway, Suite 100
Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com

060342BR v1.6