

Performance Accelerated

Mellanox InfiniBand Adapters Provide Advanced
Data Center Performance, Efficiency and Scalability



VPI / InfiniBand

Mellanox InfiniBand Host Channel Adapters (HCA) enable the highest data center performance through the delivery of state-of-the-art solutions for High-Performance Computing (HPC), Machine Learning, Data Analytics, Database, Cloud and Storage Platforms.

Today's exponential data growth is driving the need for intelligent and faster interconnect.

Leveraging faster speeds and innovative In-Network Computing technologies, Mellanox InfiniBand adapters enable clustered databases, parallelized applications, transactional services, and high-performance embedded I/O applications to achieve significant performance improvements and scale, lowering cost per operation and increasing overall ROI.

Mellanox delivers the most technologically advanced HCAs. Providing best-in-class performance and efficiency, they are the ideal solution for HPC clusters that demand high bandwidth, high message rate and low latency to achieve the highest server efficiency and application productivity.

With RDMA traffic consolidation and hardware acceleration for virtualization, Mellanox HCAs provide optimal I/O services such as high bandwidth and server utilization to achieve the maximum return on investment (ROI) for data centers, high scale storage systems and cloud computing.

By providing Virtual Protocol Interconnect® (VPI), Mellanox HCAs offer the flexibility of connectivity for InfiniBand and Ethernet protocols within the same adapter.

World-Class Performance and Scale

Mellanox InfiniBand adapters deliver industry-leading bandwidth with ultra low-latency and efficient computing for performance-driven server and storage clustering applications. Network protocol processing and data movement such as RDMA and Send/Receive semantics are completed in the adapter without CPU intervention. Application acceleration, participation in the Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)[™], and GPU communication acceleration bring further levels of performance improvement. The innovative acceleration technology in Mellanox InfiniBand adapters enables higher cluster efficiency and large scalability to hundreds of thousands of nodes.

Complete End-to-End HDR InfiniBand Networking

ConnectX adapters are part of Mellanox's full HDR 200Gb/s InfiniBand end-to-end portfolio for data centers and high-performance computing systems, which includes switches, application acceleration packages, and cables. Mellanox's Quantum[™] family of HDR InfiniBand switches and Unified Fabric Management software incorporate advanced tools that simplify networking management and installation, and provide the needed capabilities for the highest scalability and future growth. Mellanox's HPC-X[™] collectives, messaging, and storage acceleration packages deliver additional capabilities for the ultimate server performance, and the line of HDR copper and fiber cables ensure the highest interconnect performance. With Mellanox end to end, IT managers can be assured of the highest performance and most efficient network fabric.

BENEFITS

- World-class cluster performance
- Networking and storage access
- Efficient use of compute resources
- Guaranteed bandwidth and low-latency services
- Smart interconnect for x86, Power, Arm, and GPU-based compute and storage platforms
- Increased VM per server ratio
- Virtualization acceleration
- Scalability to hundreds-of-thousands of nodes
- Open Data Center Committee (ODCC) compatible

TARGET APPLICATIONS

- High-performance parallel computing
- Machine Learning and data analysis platforms
- Clustered database applications and high-throughput data warehousing
- Latency-sensitive applications such as financial analysis and high frequency trading
- Embedded systems leveraging high performance internal performance
- Performance storage applications such as backup, restore, mirroring, etc.

Virtual Protocol Interconnect® (VPI)

Mellanox's VPI flexibility enables any standard networking, clustering, storage, and management protocol to seamlessly operate over any converged network leveraging a consolidated software stack. Each port can operate on InfiniBand or Ethernet fabrics, and supports IP over InfiniBand (IPoIB) and RDMA over Converged Ethernet (RoCE). VPI simplifies I/O system design and makes it easier for IT managers to deploy infrastructure that meets the challenges of a dynamic data center.

I/O Virtualization

Mellanox adapters provide comprehensive support for virtualized data centers with Single-Root I/O Virtualization (SR-IOV) allowing dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server. I/O virtualization on InfiniBand gives data center managers better server utilization and LAN and SAN unification while reducing cost, power, and cable complexity.

Mellanox Multi-Host Solution

Mellanox's Multi-Host® technology provides high flexibility and major savings in building next generation, scalable high-performance data centers. Mellanox Multi-Host connects multiple compute or storage hosts into a single interconnect adapter, separating the adapter PCIe interface into multiple and independent PCIe interfaces with no performance degradation. The technology enables designing and building new scale-out heterogeneous compute and storage racks with direct connectivity between compute elements, storage elements and the network, better power and performance management, while achieving maximum data processing and data transfer at minimum capital and operational expenses.

Various Form Factors

Mellanox adapter cards are available in a variety of form factors to meet every data center's specific needs.

- Open Compute Project (OCP) cards integrate into the most cost-efficient, energy-efficient and scalable enterprise and hyperscale data centers, delivering leading connectivity for performance-driven server and storage applications. The OCP Mezzanine adapter form factor is designed to mate into OCP servers.
- Mellanox Socket Direct® cards split the PCIe bus into 2 buses, such that each CPU socket gets direct connectivity to the network. With this direct connectivity, traffic can bypass the inter-processors interface, optimizing performance and reducing latency for dual socket servers. Also, each CPU handles only its own traffic, improving CPU utilization. GPUDirect® RDMA is also enabled for all CPU/GPU pairs, ensuring that all GPUs are linked to those CPUs that are closest to the adapter card.
- Mellanox Socket Direct cards enable HDR 200Gb/s transmission rates for PCIe Gen3 servers, leveraging two PCIe Gen3 x16 slots. Other flavors of Socket Direct cards split 16-lane PCIe into two 8-lane buses.

Accelerated Storage

A consolidated compute and storage network provides significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols leveraging InfiniBand RDMA result in high-performance storage access. Adapters support SRP, iSER, NFS RDMA, SMB Direct, SCSI and iSCSI, as well as NVMe over Fabrics storage protocols. ConnectX adapters also offer a flexible Signature Handover mechanism based on the advanced T-10/DIF implementation.

Enabling High Performance Computing (HPC) Applications

Mellanox InfiniBand/VPI adapters are the perfect solution for the evolving data-centric paradigm. Technologies within this model include the innovative In-Network Computing offloads that transform the data center interconnect into a “distributed CPU,” and “distributed memory,” overcoming performance bottlenecks and enabling faster and more scalable data analysis.

Mellanox’s advanced In-Network Computing accelerations and RDMA offload capabilities optimize the performance of a wide variety of HPC and machine learning systems in bioscience, media, automotive design, CFD and manufacturing, weather research, oil and gas, and other markets. As a core In-Networking Computing technology, Mellanox Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)[™] optimizes MPI operations performance, decreasing data load on the network and dramatically reducing MPI operations time, while freeing up CPU resources needed for other tasks.

Software Support

All Mellanox adapters are supported by a full suite of drivers for Microsoft Windows, Linux and FreeBSD major distributions. The adapters support OpenFabrics-based RDMA protocols and software and are compatible with configuration and management tools from various OEMs and operating system vendors.

ConnectX[®]-6

ConnectX-6 is the world’s first 200Gb/s HDR InfiniBand and Ethernet network adapter card, offering world-leading performance, smart offloads and in-network computing. ConnectX-6 with VPI provides two ports of 200Gb/s supporting HDR, HDR100, EDR, FDR, QDR, DDR and SDR InfiniBand speeds, as well as 200, 100, 50, 40, 25, 10, and 1 Gb/s Ethernet speeds. ConnectX-6 supports sub-600 nanosecond latency, up to 215 million messages/sec, an embedded PCIe switch, and NVMe over Fabric. In addition to all the features included in earlier versions of ConnectX, ConnectX-6 offers Mellanox Multi-Host support for up to 8 hosts and block-level encryption as a crucial innovation to network security, altogether delivering the highest performance, most secure and extremely flexible solution for today’s demanding applications and markets.

ConnectX[®]-5


Intelligent ConnectX-5 adapter cards support co-design and in-network computing, while introducing acceleration engines for maximizing HPC, data analytics and storage platforms. Also supports two ports of EDR 100Gb/s InfiniBand and 100Gb/s Ethernet connectivity, sub-600 nanosecond latency, a very high message rate, plus PCIe switch and NVMe over Fabric offloads. Includes new Message Passing Interface (MPI) offloads, e.g., MPI Tag Matching and MPI AlltoAll operations, advanced dynamic routing, and new data algorithms capabilities. Offers advanced application offloads supporting 100Gb/s for servers without x16 PCIe slots.

ConnectX[®]-4

Mellanox ConnectX-4 adapter cards with VPI combine the flexibility of InfiniBand and Ethernet protocol connectivity within the same adapter, to support EDR 100Gb/s InfiniBand and 100Gb/s Ethernet connectivity. Enabling extremely high throughput and low latency, ConnectX-4 is a high performance and flexible solution for data analytics, Web access and storage platforms. Enabling efficient I/O consolidation, the ConnectX-4 adapter card significantly reduces data center costs and complexity.

ConnectX[®]-3 Pro

Mellanox’s ConnectX-3 Pro Virtual Protocol Interconnect (VPI) adapter delivers high throughput across the PCI Express 3.0 host bus, by providing a FDR 56Gb/s InfiniBand and 40Gb Ethernet interconnect solution (up to 56GbE when connected to a Mellanox switch). Enabling fast transaction latency (less than 1usec), and delivery of more than 90M MPI messages/second, makes ConnectX[®]-3 Pro a highly-scalable, suitable solution for transaction-demanding applications.




General Specs


Ports	Single, Dual	Single, Dual	Single, Dual	Single, Dual
Port Speed (Gb/s)	IB: SDR, DDR, QDR, FDR10, FDR Eth: 1, 10, 40, 56	IB: SDR, DDR, QDR, FDR10, FDR, EDR Eth: 1, 10, 25, 40, 50, 56, 100	IB: SDR, DDR, QDR, FDR, EDR Eth: 1, 10, 25, 40, 50, 100	IB: SDR, DDR, QDR, FDR, EDR, HDR100, HDR Eth: 1, 10, 25, 40, 50, 100, 200
PCI Express (PCIe)	Gen3 x8	Gen3 x8 Gen3 x16	Gen3 x16 Gen4 x16	Gen3/4 x16 32 lanes as 2x Gen3 x 16-lane PCIe
Connectors	QSFP+	QSFP28	QSFP28	QSFP56
RDMA Message Rate (million msgs/sec)	36	150	200 (ConnectX-5 Ex, Gen4 server) 165 (ConnectX-5, Gen3 server)	215
Latency (us)	0.64	0.6	0.6	0.6
Typical Power (2 ports, max. speed)	6.2W	16.3W	19.3W (ConnectX-5 Ex, Gen4 server) 16.2W (ConnectX-5, Gen3 server)	Contact Mellanox Support
RDMA	✓	✓	✓	✓
OOO RDMA (Adaptive Routing)	—	—	✓	✓
Dynamically Connected Transport	—	✓	✓	✓
Mellanox Multi-Host	—	4 hosts	4 hosts	8 hosts

Storage


NVMe-oF Target Offload	—	—	✓	✓
T-10 Dif/Signature Handover	—	✓	✓	✓




ConnectX-3^{*}
PRO



ConnectX-4^{*}



ConnectX-5



ConnectX-6

Virtualization

	127 VFs	16 PFs per port, 254 VFs	16 PFs per port, 1000 VFs per port	16 PFs per port, 1K VFs per port
SR-IOV				
Congestion Control (QCN, ECN)	✓	✓	✓	✓
MPI Tag Matching Offload	—	—	✓	✓
OVS Offload	—	✓	✓	✓
VM Isolation and Protection	—	—	✓	✓

Security

Block-level XTS-AES hardware encryption	—	—	—	✓
FIPS Capable	—	—	—	✓

Management

Hairpin (Host Chaining)	—	—	✓	✓
Host Management	—	✓	✓	✓
Multi-Host Isolation and Protection	—	—	✓	✓

QoS

Packet Pacing	—	✓	✓	✓
---------------	---	---	---	---

Form Factors

OCP	✓	—	✓	✓
Mellanox Socket Direct	—	—	✓	✓

*ConnectX-5 and ConnectX-6 offer richer feature sets that are recommended for the latest market applications.

World-leading HPC centers are making the smart decision and choosing **InfiniBand**.

This is why:

Julich Supercomputer Centre



The Julich Supercomputer Centre chose HPC Testimonia for a balanced, Co-Design approach to its interconnect, providing low latency, high throughput, and future scalability to its cluster, which contributes to projects in the areas of energy, environment, and brain research.



“ We chose a co-design approach, the appropriate hardware, and designed the system. This system was of course targeted at supporting in the best possible manner our key applications. The only interconnect that really could deliver that was HPC Testimonia. ”

CHPC South Africa



The Centre for High Performance Computing in South Africa, the largest HPC facility in Africa, chose HPC Testimonia to enhance and unlock the vast potential of its system, which provides high end computational resources to a broad range of users in fields such as bioinformatics, climate research, material sciences, and astronomy.



“ The heartbeat of the cluster is the interconnect. Everything is about how all these processes shake hands and do their work. InfiniBand and the interconnect is, in my opinion, what defines HPC. ”



VPI / InfiniBand

For detailed information on features, compliance, and compatibility, please see each product's specific product brief.



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085

Tel: 408-970-3400 • Fax: 408-970-3403

www.mellanox.com

This brochure describes hardware features and capabilities. Please refer to the driver release notes on mellanox.com for feature availability. Product images may not include heat sync assembly; actual product may differ.

© Copyright 2020. Mellanox Technologies. All rights reserved.

Mellanox, Mellanox logo, ConnectX, GPUDirect, Mellanox Multi-Host, Mellanox Socket Direct, HPC-X, UFM, and Virtual Protocol Interconnect are registered trademarks of Mellanox Technologies, Ltd. Scalable Hierarchical Aggregation and Reduction Protocol (SHARP) and Mellanox Quantum are trademarks of Mellanox Technologies, Ltd. All other trademarks are property of their respective owners.

3525BR
Rev 11.7