**Mellanox** ®
TECHNOLOGIES

# ConnectX®

## Single/Dual- Port Adapter Cards supporting up to 40Gb/s InfiniBand

ConnectX adapter cards provide the highest performing and most flexible interconnect solution for Enterprise Data Centers, High-Performance Computing, and Embedded environments. Clustered data bases, parallelized applications, transactional services and high-performance embedded I/O applications will achieve significant performance improvements resulting in reduced completion time and lower cost per operation.

*Single- and Dual-Port QSFP*

ConnectX based adapter cards simplifies network deployment by consolidating cables and enhancing performance in virtualized server environments.

### World-Class Performance Over InfiniBand

ConnectX delivers low-latency and high-bandwidth for performance-driven server and storage clustering applications. These applications will benefit from the reliable transport connections and advanced multicast support offered by ConnectX. Network protocol processing and data movement overhead such as InfiniBand RDMA and Send/Receive semantics are completed in the adapter without CPU intervention. Servers supporting PCI Express 2.0 with 5GT/s will be able to take advantage of 40Gb/s InfiniBand, balancing the I/O requirement of these high-end servers.

### TCP/UDP/IP Acceleration

Applications utilizing TCP/UDP/IP transport can achieve industry-leading throughput over InfiniBand. The hardware-based stateless offload engines in ConnectX reduce the CPU overhead of IP packet transport, freeing more processor cycles to work on the application.

### I/O Virtualization

ConnectX support for hardware-based I/O virtualization provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server. I/O virtualization with ConnectX gives data center managers better server utilization and LAN and SAN unification while reducing cost, power, and cable complexity.

### Storage Accelerated

A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols leveraging InfiniBand RDMA result in high-performance storage access. Fibre Channel frame encapsulation (FCoIB or FCoE) and hardware offloads enable simple connectivity to Fibre Channel SANs.

### Software Support

All Mellanox adapter cards are compatible with TCP/IP and OpenFabrics-based RDMA protocols and software. They are also compatible with InfiniBand and cluster management software available from OEMs. The adapter cards are compatible with major operating system distributions.
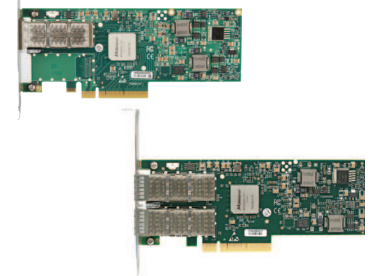
## HIGHLIGHTS

– Industry-standard technology
– World-class cluster performance
– High-performance networking and storage access
– Guaranteed bandwidth and low-latency services
– Reliable transport
– I/O consolidation
– Virtualization acceleration
– Scales to tens-of-thousands of nodes
– 1us MPI ping latency
– Selectable 10, 20, or 40Gb/s InfiniBand or 10GigE per port
– Single- and Dual-Port options available
– PCI Express 2.0 (up to 5GT/s)
– CPU offload of transport operations
– End-to-end QoS and congestion control
– Hardware-based I/O virtualization
– TCP/UDP/IP stateless offload
– Fibre Channel encapsulation (FCoIB or FCoE)

## SPECIFICATIONS

### INFINIBAND
- IBTA Specification 1.2.1 compliant
- 10, 20, or 40Gb/s per port
- RDMA, Send/Receive semantics
- Hardware-based congestion control
- Atomic operations
- 16 million I/O channels
- 256 to 4Kbyte MTU, 1Gbyte messages
- 9 virtual lanes: 8 data + 1 management

### ENHANCED INFINIBAND
- Hardware-based reliable transport
- Hardware-based reliable multicast
- Extended Reliable Connected transport
- Enhanced Atomic operations
- Fine grained end-to-end QoS

### HARDWARE-BASED I/O VIRTUALIZATION
- Single Root IOV
- Address translation and protection
- Multiple queues per virtual machine
- VMware NetQueue support
- PCISIG IOV compliant

### ADDITIONAL CPU OFFLOADS
- TCP/UDP/IP stateless offload
- Intelligent interrupt coalescence
- Compliant to Microsoft RSS and NetDMA

### STORAGE SUPPORT
- T10-compliant Data Integrity Field support
- Fibre Channel over InfiniBand or Ethernet

### CONNECTIVITY
- Interoperable with InfiniBand or 10GigE switches
- microGiGaCN or QSFP connectors
- 20m+ (10Gb/s), 10m+ (20Gb/s) or 7m+ (40Gb/s) of passive copper cable
- External optical media adapter and active cable support

### MANAGEMENT AND TOOLS INFINIBAND
- OpenSM
- Interoperable with third-party subnet managers
- Firmware and debug tools (MFT, IBDIAG)

### OPERATING SYSTEMS/DISTRIBUTIONS
- Novell SLES, Red Hat Enterprise Linux (RHEL), Fedora, and other Linux distributions
- Microsoft Windows Server 2003/2008/CCS 2003
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF)
- VMware ESX Server 3.5, Citrix XenServer 4.1

### PROTOCOL SUPPORT
- Open MPI, OSU MVAPICH, HP MPI, Intel MPI, MS MPI, Scali MPI
- TCP/UDP, IPoIB, SDP, RDS
- SRP, iSER, NFS RDMA, FCoIB, FCoE
- uDAPL

### Ordering Information

| Ordering Part Number | Ports | Host Bus | Power (Typ) | Dimensions w/o Brackets |
|---|---|---|---|---|
| MHGH28-XTC | Dual 4X 20Gb/s InfiniBand | PCIe 2.0 2.5GT/s | 11.0W | 13.6cm x 6.4cm |
| MHGH29-XTC | Dual 4X 20Gb/s InfiniBand | PCIe 2.0 5.0GT/s | 11.6W | 13.6cm x 6.4cm |
| MHJH29-XTC | Dual 4X 40Gb/s InfiniBand | PCIe 2.0 5.0GT/s | 12.2W | 13.6cm x 6.4cm |
| MHRH29-XTC | Dual 4X QSFP 20Gb/s InfiniBand | PCIe 2.0 5.0GT/s | 11.6W | 16.8cm x 6.4cm |
| MHQH29-XTC | Dual 4X QSFP 40Gb/s InfiniBand | PCIe 2.0 5.0GT/s | 12.2W | 16.8cm x 6.4cm |

Single-port options also available

**Mellanox** TECHNOLOGIES

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