



ADAPTER CARDS

ConnectX®

Dual-Port 20 and 40Gb/s InfiniBand Adapters

ConnectX® 20 and 40Gb/s InfiniBand dual-port adapters deliver low-latency and high-bandwidth for performance-driven server and storage clustering applications in Enterprise Data Center and High-Performance Computing environments. Clustered data bases, parallelized applications and transactional services applications will achieve significant performance improvements resulting in reduced completion time and lower cost per operation. ConnectX simplifies network deployment by consolidating clustering, communications, storage, and management I/O and by providing enhanced performance in virtualized server environments.

World-Class Performance and Scalability

Clustered applications running on multi-socket servers using multi-core processors will benefit from the reliable transport connections and advanced multicast support offered by ConnectX. End-to-end Quality of Service (QoS) enables partitioning and guaranteed service levels while hardware-based congestion control prevents hot spots from degrading the effective throughput. ConnectX is capable of scaling to tens-of-thousands of server and storage nodes.

Hardware Offload Architecture

Clustered and client/server applications achieve maximum performance over ConnectX because CPU cycles are available to focus on critical application processing instead of networking functions. Network protocol processing and data movement overhead such as RDMA and Send/Receive semantics are completed in the adapter without CPU intervention. Applications utilizing TCP/UDP/IP transport can achieve industry-leading 40Gb/s throughput when run over ConnectX and its hardware-based stateless offload engines.

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 unified InfiniBand cluster for computing and storage 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 over InfiniBand (FCoIB) hardware offloads enable simple connectivity to Fibre Channel SANs.

Software Support

All Mellanox adapter cards are compatible with legacy 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 supported with major operating system distributions.



BENEFITS

- World-class cluster performance
- High-performance networking and storage access
- Guaranteed bandwidth and low-latency services
- Reliable transport
- End-to-end storage integrity
- I/O consolidation
- Virtualization acceleration
- Scales to tens-of-thousands of nodes

KEY FEATURES

- 1us MPI ping latency
- 20 or 40Gb/s InfiniBand ports
- CPU offload of transport operations
- End-to-end QoS and congestion control
- Hardware-based I/O virtualization
- TCP/UDP/IP stateless offload



FEATURE SUMMARY**INFINIBAND**

- IBTA Specification 1.2 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
- 1GB 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 compliance

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 (FCoIB)

COMPATIBILITY**PCI EXPRESS INTERFACE**

- PCIe Base 2.0 compliant, 1.1 compatible
- 2.5GT/s or 5.0GT/s link rate x8 (20+20Gb/s or 40+40Gb/s bidirectional bandwidth)
- Support for MSI/MSI-X mechanisms

MANAGEMENT AND TOOLS

- OpenSM
- Interoperable with third-party subnet managers
- Firmware and debug tools (MFT, IBADM)

OPERATING SYSTEMS/DISTRIBUTIONS

- Novell SLES, Red Hat Enterprise Linux (RHEL), Fedora, and other Linux distributions
- Microsoft Windows Server 2007/2008/CCS

PROTOCOL SUPPORT

- Open MPI, OSU MVAPICH, HP MPI, Intel MPI, MS MPI, Scali MPI
- IPoIB, SDP, RDS
- SRP, iSER, FCoIB and NFS RDMA
- uDAPL

COMPLIANCE**SAFETY**

- US/Canada: cTUVus
- EU: IEC60950
- International: CB

EMC (EMISSIONS)

- USA: FCC, Class A
- Canada: ICES, Class A
- EU: CE Mark (EN55022 Class A, EN55024, EN61000-3-2, EN61000-3-3)
- Japan: VCCI, Class A
- Korea: KCC Class A
- Australia/New Zealand: C-Tick Class A

ENVIRONMENTAL

- EU: IEC 60068-2-64: Random Vibration
- EU: IEC 60068-2-29: Shocks, Type I / II
- EU: IEC 60068-2-32: Fall Test

OPERATING CONDITIONS

- Operating temperature: 0 to 55° C
- Air flow: 200LFM @ 55° C
- Requires 3.3V, 12V supplies

SPECIFICATIONS

- Dual 4X InfiniBand ports
- PCI Express 2.0 x8 (1.1 compatible)
- Single chip architecture
- Low profile PCIe adapter card
- RoHS-R5 compliant
- 37-month warranty from date of manufacturing

**Adapter Cards**

Ordering Part Number	Ports	Host Bus	Power (Typ)
44R8728	Dual 4X 20Gb/s	PCIe 2.0 5.0GT/s	11.6W
46M2203	Single QSFP 40Gb/s	PCIe 2.0 5.0GT/s	12.6W
46M2199	Dual QSFP 40Gb/s	PCIe 2.0 5.0GT/s	13.5W



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

© Copyright 2009, Mellanox Technologies. All rights reserved.
 Mellanox, ConnectX, InfiniBlast, InfiniBridge, InfiniHost, InfiniRISC, InfiniScale, and InfiniPCI are registered trademarks of Mellanox Technologies, Ltd. Virtual Protocol Interconnect and BridgeX are trademarks of Mellanox Technologies, Ltd. All other trademarks are property of their respective owners.