



ADAPTER CARD

ConnectX[®]-2 VPI

Dual-Port Adapter with *Virtual Protocol Interconnect*[™]
InfiniBand and Ethernet Connectivity

ConnectX-2 adapter card with *Virtual Protocol Interconnect* (VPI) supporting InfiniBand and Ethernet connectivity provides the highest performing and most flexible interconnect solution for Enterprise Data Centers, High-Performance Computing, and Embedded environments. InfiniBand and 10 Gigabit Ethernet provides the highest performing interconnect for clustered data bases, parallelized applications, transactional services and high-performance embedded I/O applications and seamless client-server connectivity over Ethernet.

Virtual Protocol Interconnect

VPI-enabled adapters make it possible for any standard networking, clustering, storage, and management protocol to seamlessly operate over any converged network leveraging a consolidated software stack. The QSFP port supports up to 40Gb/s InfiniBand while the SFP+ port supports 10 Gigabit Ethernet. Both ports can utilize passive copper, active copper, or optical cables. ConnectX-2 with 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.

World-Class Performance

InfiniBand — ConnectX-2 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-2. Network protocol processing and data movement overhead such as InfiniBand RDMA and Send/Receive semantics are completed in the adapter without CPU intervention improving the server's overall efficiency. ConnectX-2 enables large scalability to tens of thousands of nodes.

Data Center Bridging — ConnectX-2 delivers similar low-latency and high-bandwidth performance leveraging Ethernet with DCB support. Low Latency Ethernet (LLE) provides efficient RDMA transport over Layer 2 Ethernet utilizing DCB's enhancements to IEEE 802.1 bridging. The LLE software stack maintains existing and future compatibility for bandwidth and latency sensitive clustering applications. With link-level interoperability in existing Ethernet infrastructure, Network Administrators can leverage existing data center fabric management solutions.

TCP/UDP/IP Acceleration — Applications utilizing TCP/UDP/IP transport can achieve industry-leading throughput over InfiniBand or 10 Gigabit Ethernet. The hardware-based stateless offload engines in ConnectX-2 reduce the CPU overhead of IP packet transport, freeing more processor cycles to work on the application.

I/O Virtualization — ConnectX-2 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-2 gives data center managers better server utilization and LAN and SAN unification while reducing cost, power, and complexity.

Storage Accelerated — A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access



ConnectX-2 VPI Adapter Card



BENEFITS

- One adapter for InfiniBand, 10 Gigabit Ethernet and Data Center Bridging fabrics
- 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

KEY FEATURES

- Virtual Protocol Interconnect
- 1us MPI ping latency over InfiniBand
- One 10, 20, or 40Gb/s InfiniBand
- One 10 Gigabit Ethernet SFP+ port
- 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 (FCoB or FCoE)
- RoHS-R6

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.



FEATURE SUMMARY

INFINIBAND

- IBTA Specification 1.2.1 compliant
- 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

ETHERNET

- IEEE 802.3ae 10 Gigabit Ethernet
- IEEE 802.3ad Link Aggregation and Failover
- IEEE 802.1Q, .1p VLAN tags and priority
- IEEE P802.1au D2.0 Congestion Notification
- IEEE P802.1az D0.2 ETS
- IEEE P802.1bb D1.0 Priority-based Flow Control
- Multicast
- Jumbo frame support (10KB)
- 128 MAC/VLAN addresses per port

HARDWARE-BASED I/O VIRTUALIZATION

- Single Root IOV
- Address translation and protection
- Dedicated adapter resources
- Multiple queues per virtual machine
- Enhanced QoS for vNICs
- VMware NetQueue support

ADDITIONAL CPU OFFLOADS

- TCP/UDP/IP stateless offload
- Intelligent interrupt coalescence

STORAGE SUPPORT

- Fibre Channel over InfiniBand or Ethernet
- T11-compliant frame format

COMPATIBILITY

PCI EXPRESS INTERFACE

- PCIe Base 2.0 compliant, 1.1 compatible
- 5.0GT/s link rate x8 40+40Gb/s bidirectional bandwidth)
- Auto-negotiates to x8, x4, x2, or x1
- Support for MSI/MSI-X mechanisms

CONNECTIVITY

- Interoperable with InfiniBand and 10GigE switches

InfiniBand

- 20m+ (10Gb/s), 10m+ (20Gb/s) or 7m+ (40Gb/s) of pasive copper cable
- 300m of optical cable

Ethernet

- 10m+ direct attached copper cable
- 100m (OM-2) or 300m (OM-3) of multimode fiber cable, duplex LC connector

MANAGEMENT AND TOOLS

InfiniBand

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

Ethernet

- MIB, MIB-II, MIB-II Extensions, RMON, RMON 2
- Configuration and diagnostic tools

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

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

Adapter Card

Ordering Part Number	Ports	Host Bus	Power (Typ)	Dimensions w/o Brackets
MHZH29-XTR	4X QSFP 40Gb/s InfiniBand , SFP+ 10 Gigabit Ethernet	PCIe 2.0 5.0GT/s	TBD	TBD



350 Oakmead Parkway, 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, PhyX and BridgeX are trademarks of Mellanox Technologies, Ltd. All other trademarks are property of their respective owners.