

ConnectX[®]

Single/Dual-Port Adapter Devices with *Virtual Protocol Interconnect*[®]

ConnectX adapter devices with *Virtual Protocol Interconnect* (VPI) provide the highest-performing and most flexible interconnect solution for Blade Server and Landed on Motherboard designs used in Enterprise Data Centers, High-Performance Computing, and Embedded environments.

With ConnectX, clustered data bases, parallelized applications, transactional services and high-performance embedded I/O applications will achieve significant performance improvements with reduced completion time and lower cost per operation. ConnectX with VPI also simplifies network deployment by consolidating cables.

Virtual Protocol Interconnect

VPI-enabled adapters enable any standard networking, clustering, storage, and management protocol to seamlessly operate over any converged network leveraging a consolidated software stack. Each ConnectX port can automatically identify and operate on InfiniBand, Ethernet, or Data Center Ethernet (DCE) fabrics. ConnectX

with VPI simplifies I/O system design and makes it easier for IT managers to deploy dynamic data center infrastructure.

World-Class Performance


InfiniBand — ConnectX delivers low latency and high bandwidth for performance-driven server and storage clustering applications. 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 can take advantage of 40Gb/s InfiniBand.

Data Center Ethernet — ConnectX delivers similar low-latency and high-bandwidth performance over DCE. Applications running on servers utilizing InfiniBand protocols over Ethernet (IBoE) will benefit from InfiniBand RDMA and Send/Receive accelerations and DCE's lossless fabric. Applications qualified using the OpenFabrics protocol stack can be easily deployed on IBoE.

TCP/UDP/IP Acceleration — Applications utilizing TCP/UDP/IP transport can achieve industry-leading throughput over InfiniBand or Ethernet. 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 supports hardware-based I/O virtualization, providing dedicated adapter resources and guaranteed isolation

ConnectX[®]



HIGHLIGHTS

BENEFITS

- One design for InfiniBand, 10Gig Ethernet or Data Center Ethernet applications
- 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
- Small PCB footprint

KEY FEATURES

- Virtual Protocol Interconnect
- Single chip architecture
 - *Integrated SerDes*
 - *No local memory needed*
- 1us MPI ping latency
- Selectable 10, 20, or 40Gb/s InfiniBand or 10GigE per 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 (FCoIB or FCoE)




Figure 1. ConnectX Block Diagram

and protection for virtual machines (VM) within the server. ConnectX gives data center managers better server utilization and LAN/SAN unification while reducing costs, power, and complexity.

Storage Acceleration — A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can leverage InfiniBand RDMA for high-performance storage access. Fibre Channel frame encapsulation (FCoB or FCoE) and hardware offloads enable simple connectivity to Fibre Channel SANs.

Software Support

All Mellanox adapter devices 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. Adapters based on the devices are compatible with major operating system distributions.



FEATURE SUMMARY

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
- 2GB messages
- 9 virtual lanes: 8 data + 1 management

ENHANCED INFINIBAND

- Hardware-based reliable transport
- Hardware-based reliable multicast
- Scalable Reliable Connected transport
- Enhanced Atomic operations
- Fine grained end-to-end QoS

ETHERNET

- IEEE Std 802.3ae 10 Gigabit Ethernet
- IEEE Std 802.3ak 10GBASE-CX4
- IEEE Std 802.3ap Backplanes
- IEEE Std 802.3ad Link Aggregation and Failover

- IEEE Std 802.3x Pause
- IEEE Std 802.1Q VLAN tags
- IEEE Std 802.1p Priorities
- Multicast
- Jumbo frame support (10KB)
- 128 MAC/VLAN addresses per port
- MAC and VLAN based filtering
- Class Based Flow Control / Per Priority Pause

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 or Ethernet

COMPATIBILITY

CPU

- AMD X86, X86_64
- Intel X86, EM64T, IA-32, IA-64
- SPARC
- PowerPC, MIPS, and Cell

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)
- Auto-negotiates to x8, x4, x2, or x1
- Support for MSI/MSI-X mechanisms

CONNECTIVITY

- Interoperable with InfiniBand switches
- Drives copper cables or backplanes

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 2007/2008/CCS
- 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, IPoB, SDP, RDS
- SRP, iSER, NFS RDMA, FCoB, FCoE
- uDAPL



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