



ConnectX™

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
- Small PCB footprint

KEY FEATURES

- Single chip architecture
 - Integrated SerDes
 - No local memory needed
- 1.2us MPI ping latency
- 10 or 20Gb/s InfiniBand ports
- 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

SPECIFICATIONS

- Dual 4X InfiniBand ports
- PCI Express 2.0 x8 (1.1 compatible)
- Management interfaces (DMTF compatible, Fast Management Link)
- 4x 16MB serial Flash interface
- Dual I²C interfaces
- IEEE 1149.1 boundary-scan JTAG
- Link status LED indicators
- General purpose I/O
- 21 x 21mm HFCBGA
- RoHS-5 compliant
- Requires 3.3V, 2.5V, 1.8V, 1.2V supplies


Mellanox ConnectX™ IB

Dual-Port InfiniBand Adapter Devices with PCI Express 2.0

Mellanox ConnectX IB InfiniBand Host Channel Adapter (HCA) devices deliver low-latency and high-bandwidth for performance-driven server and storage clustering applications in 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. ConnectX IB simplifies network deployment by consolidating clustering, communications, storage, and management I/O and by providing enhanced performance in virtualized server environments. The devices are well suited for Blade Server and Landed on Motherboard designs due to their small overall footprint requirement.

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



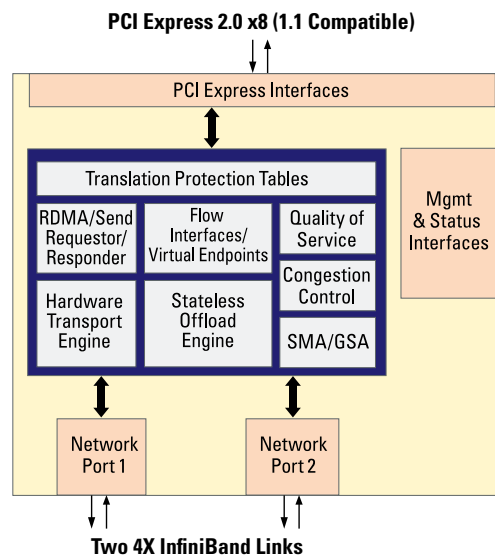
IB. Servers supporting PCI Express 2.0 with 5GT/s will be able to utilize the full potential of 20Gb/s InfiniBand, balancing the I/O requirement of these high-end servers. End-to-end Quality of Service (QoS) enables partitioning and guaranteed service levels while hardware-based congestion control prevents 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 IB 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 device without CPU intervention. Applications utilizing TCP/UDP/IP transport can achieve industry-leading throughput when run over ConnectX IB and its hardware-based stateless offload engines.

I/O Virtualization

ConnectX IB support for hardware-based I/O virtualization is complementary to Intel and AMD virtualization technologies. Virtual machines (VM) within the server are enabled with dedicated I/O adapter resources and guaranteed isolation and protection. Hypervisor offload features remove software-based virtualization overheads and free



ConnectX IB Block Diagram

ConnectX™ IB Dual-Port InfiniBand Adapter Devices with PCI Express 2.0

up CPU cycles enabling native OS performance for VMs and higher server utilization by supporting more VMs per physical server.

Storage Accelerated

A unified InfiniBand network 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. Data reliability is improved through the use of T10-compliant Data Integrity Field (DIF). Fibre Channel (FC) over InfiniBand (FCoIB) features enable the use of cost-effective bridges for connecting to FC SANs.

Software Support

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

Mellanox Advantage

Mellanox is the leading supplier of industry standard InfiniBand HCAs and switch silicon. Our products have been deployed in clusters scaling to thousands of nodes and are being deployed end-to-end in data centers and Top500 systems around the world.

FEATURE SUMMARY

INFINIBAND

- IBTA Specification 1.2 compliant
- 10 or 20Gb/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
- Service oriented I/O
- Fine grained end-to-end QoS

HARDWARE-BASED I/O VIRTUALIZATION

- Address translation and protection
- Multiple queues per virtual machine
- Native OS performance
- Complimentary to Intel and AMD I/OMMU

ADDITIONAL CPU OFFLOADS

- TCP/UDP/IP stateless offload
- Intelligent interrupt coalescence
- Full support for Intel I/OAT
- Compliant to Microsoft RSS and NetDMA

STORAGE SUPPORT

- T10-compliant Data Integrity Field support
- Fibre Channel over InfiniBand (FCoIB)

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

- 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/CCS/XP, Longhorn
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinIB)

PROTOCOL SUPPORT

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

ADAPTER SILICON

Ordering Part Number	InfiniBand Ports	Host Bus	Power (2 Ports, Typ.)
MT25408A0-FCC-SI	Dual 4X (10Gb/s)	PCIe 2.0 2.5GT/s	9.6W
MT25408A0-FCC-DI	Dual 4X (10Gb/s, 20Gb/s)	PCIe 2.0 2.5GT/s	10.1W
MT25408A0-FCC-GI	Dual 4X (10Gb/s, 20Gb/s)	PCIe 2.0 5.0GT/s	10.7W



2900 Stender Way, Santa Clara, CA 95054
 Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com