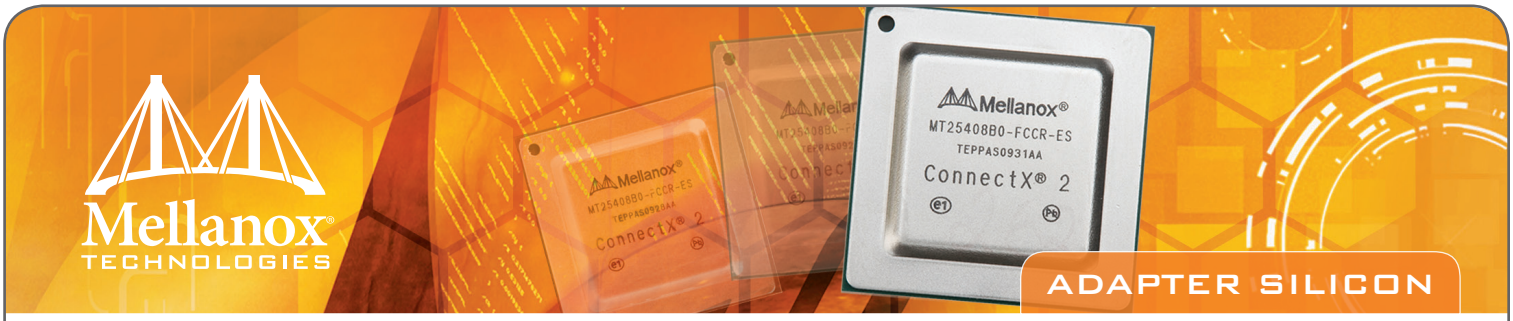

ConnectX®-2

Single/Dual-Port Adapter Device with *Virtual Protocol Interconnect*™

ConnectX-2 adapter devices with *Virtual Protocol Interconnect* (VPI) supporting InfiniBand and Ethernet connectivity 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. Clustered data bases, parallel processing, 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-2 with VPI also simplifies network deployment by consolidating cables and enhancing performance in virtualized server environments.

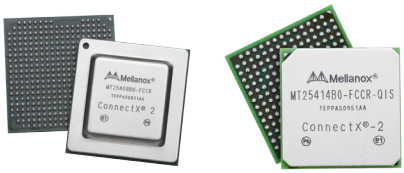

Virtual Protocol Interconnect

VPI-supported adapters enable any standard networking, clustering, storage, and management protocol to seamlessly operate over any converged network leveraging a consolidated software stack. With auto-sense capability, each ConnectX-2 port can identify and operate on InfiniBand, Ethernet, or Data Center Bridging (DCB) fabrics. FlexBoot™ superscript provides additional flexibility by enabling servers to boot from remote InfiniBand or LAN storage targets. ConnectX-2 with VPI and FlexBoot simplifies I/O system design and makes it easier for IT managers to deploy dynamic data center infrastructure.

World-Class Performance

InfiniBand — ConnectX-2 delivers low latency, high bandwidth, and computing efficiency for performance-driven server and storage clustering applications. Efficient computing is achieved by offloading from the CPU routine activities which allows more processor power for the application. Network protocol processing and data movement overhead such as InfiniBand RDMA and Send/Receive semantics are completed in the adapter without CPU intervention. CORE-Direct™ brings the next level of performance improvement by offloading application overhead (e.g. MPI collectives operations), such as data broadcasting and gathering as well as global synchronization communication routines. GPU-Direct™ provides additional efficiencies by eliminating unnecessary internal data copies to significantly reduce application run time. ConnectX-2 advanced acceleration technology enables higher cluster efficiency and 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. RDMA over Ethernet (RoE), utilizing IBTA RDMA over Converged Ethernet (RoCE) technology, provides efficient low latency RDMA transport over Layer 2

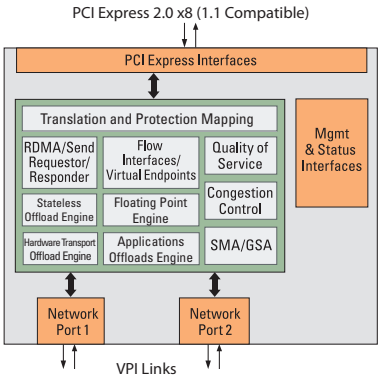

Dual-Port ConnectX-2 Block Diagram

ConnectX®-2

BENEFITS

- One design for InfiniBand, 10Gig Ethernet or 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
- 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
- CORE-Direct application offload
- GPU-Direct communication acceleration
- End-to-end QoS and congestion control
- Hardware-based I/O virtualization
- T11.3 FC-BB-5 FCoE
- 16mm X 16mm (single port)
- 21mm X 21mm (dual port)

Ethernet. The RoE software stack maintains existing and future compatibility with bandwidth and latency sensitive 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 10GigE. 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 with Virtual Intelligent Queuing (Virtual-IQ) technology 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 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. T11 compliant encapsulation (FCoIB or FCoE) with full hardware offloads simplifies the storage network while keeping existing Fibre Channel targets.

Software Support

All Mellanox adapters are supported by a full suite of drivers for Microsoft Windows, Linux distributions, VMware, and Citrix XENServer. ConnectX-2 VPI adapters support OpenFabrics-based RDMA protocols and software. Stateless offloads are fully interoperable with standard TCP/UDP/IP stacks. ConnectX-2 VPI adapters are compatible with configuration and management tools from OEMs and operating system vendors.

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

ENHANCED INFINIBAND
- Hardware-based reliable transport
- Collective operations offloads
- GPU communication acceleration
- Hardware-based reliable multicast
- Extended Reliable Connected transport
- Enhanced Atomic operations

ETHERNET
- IEEE Std 802.3ae 10 Gigabit Ethernet
- IEEE Std 802.3ap Backplanes, including FEC
- IEEE Std 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 (9.6KB)
- 128 MAC/VLAN addresses per port
- MAC and VLAN based filtering

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
- RDMA over Converged Ethernet
- TCP/UDP/IP stateless offload
- Intelligent interrupt coalescence

STORAGE SUPPORT
- T11.3 FC-BB-5 FCoE

FLEXBOOT™ TECHNOLOGY
- Remote boot over InfiniBand
- Remote boot over Ethernet
- Remote boot over iSCSI

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 and 10GigE switches
- Drives copper cables or backplanes

MANAGEMENT AND TOOLS

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

Ethernet
- MIB, MIB-II, MIB-II Extensions, RMON, RMON2
- 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/vSphere 4.0

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

OPENFABRICS SOFTWARE

INFINIBAND TRADE ASSOCIATION CERTIFIED™

Mellanox TECHNOLOGIES

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