

ConnectX[®]-3 VPI



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

ConnectX-3 adapter devices with *Virtual Protocol Interconnect (VPI)* supporting InfiniBand and Ethernet connectivity provide the highest-performing and most flexible interconnect solution for PCI Express Gen3 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-3 with VPI also simplifies system development by serving multiple fabrics with one hardware design.

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-3 port can identify and operate on InfiniBand, Ethernet, or Data Center Bridging

(DCB) fabrics. FlexBoot™ provides additional flexibility by enabling servers to boot from remote InfiniBand or LAN storage targets. ConnectX-3 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-3 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 network protocol processing and data movement overhead such as RDMA and Send/Receive semantics allowing more processor power for the application. CORE-Direct™ brings the next level of performance improvement by offloading application overhead such as data broadcasting and gathering as well as global synchronization communication routines. GPU communication acceleration provides additional efficiencies by eliminating unnecessary internal data copies to significantly reduce application run time. ConnectX-3 advanced acceleration technology enables higher cluster efficiency and large scalability to tens of thousands of nodes.

RDMA over Converged Ethernet—

ConnectX-3 utilizing IBTA RoCE technology delivers similar low-latency and high-performance over Ethernet networks. Leveraging Data Center Bridging capabilities, RoCE provides efficient low latency RDMA services over Layer



HIGHLIGHTS

BENEFITS

- One design for InfiniBand, 10/40/56Gig Ethernet or Data Center Bridging fabrics
- World-class cluster, network, and storage performance
- Guaranteed bandwidth and low-latency services
- I/O consolidation
- Virtualization acceleration
- Scales to tens-of-thousands of nodes
- Small PCB footprint

KEY FEATURES

- 1us MPI ping latency
- Up to 56Gb/s InfiniBand or 40/56 Gigabit Ethernet per port
- PCI Express 3.0 (up to 8GT/s)
- CPU offload of transport operations
- Application offload
- GPU communication acceleration
- Precision Clock Synchronization
- End-to-end QoS and congestion control
- Hardware-based I/O virtualization
- Dynamic power management
- Fibre Channel encapsulation (FCoIB or FCoE)
- Ethernet encapsulation (EoIB)
- 17mm X 17mm RoHS-R6

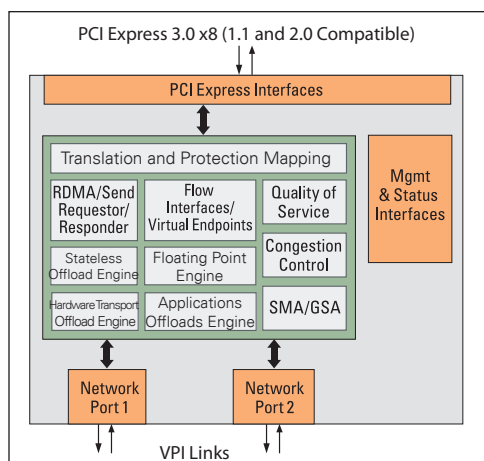


Figure 1. Dual-Port ConnectX-3 Block Diagram

2 Ethernet. With link-level interoperability in existing Ethernet infrastructure, Network Administrators can leverage existing data center fabric management solutions.

Sockets Acceleration — Applications utilizing TCP/UDP/IP transport can achieve industry-leading throughput over InfiniBand or 10/40/56GbE. The hardware-based stateless offload engines in ConnectX-3 reduce the CPU overhead of IP packet transport, Sockets acceleration software further increases performance for latency sensitive applications.

I/O Virtualization — ConnectX-3 Virtual Intelligent Queuing (Virtual-IQ) technology with SR-IOV provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server. I/O virtualization with ConnectX-3 gives data center managers better server utilization while reducing costs, power, and complexity.

Precision Data Centers — ConnectX-3 - IEEE 1588 precision time protocol circuitry synchronizes the host clock to the data center master clock for accurate data delivery time stamping and data center SLA measurements. The hardware-based mechanisms ensure high accuracy and low jitter.

Storage Acceleration — A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can utilize 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-3 VPI adapters support OpenFabrics-based RDMA protocols and software and are compatible with configuration and management tools from OEMs and operating system vendors.

FEATURE SUMMARY

INFINIBAND

- IBTA Specification 1.2.1 compliant
- Hardware-based congestion control
- 16 million I/O channels
- 256 to 4Kbyte MTU, 1Gbyte messages

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 802.3ba 40 Gigabit Ethernet
- IEEE Std 802.3ap Backplanes, including FEC
- IEEE Std 802.3ad Link Aggregation and Failover
- IEEE Std 802.3az Energy Efficient Ethernet
- IEEE Std 802.1Q, .1p VLAN tags and priority
- IEEE Std 802.1Qau Congestion Notification
- IEEE P802.1Qaz D0.2 ETS

- IEEE P802.1Qbb D1.0 Priority-based Flow Control
- IEEE Std 1588 Precision Clock Synchronization
- Jumbo frame support (9600B)

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

PCI EXPRESS INTERFACE

- PCIe Base 3.0 compliant, 1.1 and 2.0 compatible
- 2.5, 5.0, or 8.0GT/s link rate x8
- Auto-negotiates to x8, x4, x2, or x1
- Support for MSI/MSI-X mechanisms

CONNECTIVITY

- Interoperable with InfiniBand or 10/40 Ethernet switches. Interoperable with 56GbE Mellanox Switches.
- Drives copper cables or backplanes

OPERATING SYSTEMS/DISTRIBUTIONS

- Citrix XenServer 6.1
- Novell SLES, Red Hat Enterprise Linux (RHEL), Fedora, and other Linux distributions
- Microsoft Windows Server 2008/2012/2012 R2
- OpenFabrics Enterprise Distribution (OFED)
- Ubuntu 12.04
- VMware ESXi 4.x and 5.x

PROTOCOL SUPPORT

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

*This product brief describes all of the hardware features and capabilities. Please refer to the driver release notes on www.mellanox.com for feature availability.



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