INFINIBAND ADAPTER CARDS

PRODUCT BRIEF

ConnectX®-3 Pro VPI Adapters for Open Compute Project Supporting OCP Specification 2.0

Single/Dual-Port Adapters with Virtual Protocol Interconnect®

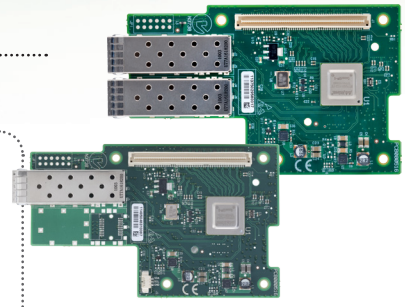

Mellanox ConnectX-3 Pro VPI adapter cards with Virtual Protocol Interconnect (VPI) supporting InfiniBand and Ethernet connectivity provide the highest performing and most flexible interconnect solution, offering high-bandwidth, low latency and industry-leading connectivity for Open Compute Project (OCP) server and storage applications in High Performance Computing, Web 2.0, Enterprise Data Centers and Cloud infrastructure.

Web2.0, public and private clouds, storage, and high performance computing are just a few applications that will benefit from the use of Mellanox ConnectX-3 Pro VPI Adapters for Open Compute Project (OCP) specification 2.0. The adapters deliver significant throughput and latency improvements resulting in faster access, real-time response and more virtual machines hosted per server. Moreover, the adapters improve network performance by increasing available bandwidth while decreasing the associated transport load on the CPU especially in virtualized server environments.

Virtual Protocol Interconnect

VPI-enabled adapters enable any standard networking, clustering and storage deployment to seamlessly operate over any converged network leveraging a consolidated software stack. With auto-sense capability, each ConnectX-3 Pro port can identify and operate on InfiniBand, Ethernet, or Data Center Bridging (DCB) fabrics. Mellanox FlexBoot™ provides additional flexibility by enabling servers to boot from remote InfiniBand or LAN storage targets. ConnectX-3 Pro with VPI and FlexBoot 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-3 Pro 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 protocol processing and data movement overhead such as RDMA and Send/Receive semantics, allowing more processor power for the application. Mellanox 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 runtime. ConnectX-3 Pro advanced acceleration technology enables higher cluster efficiency and large scalability to tens of thousands of nodes.

RDMA over Converged Ethernet –
ConnectX-3 Pro utilizing IBTA RoCE technology delivers similar low-latency and high-performance

HIGHLIGHTS

BENEFITS

- One adapter for 10/40/56 Gigabit InfiniBand, Ethernet or Data Center Bridging fabrics
- Open Compute Project Form Factor
- World-class cluster, network, and storage performance
- Guaranteed bandwidth and low-latency services
- Cutting edge performance in virtualized overlay networks (VXLAN and NVGRE)
- I/O consolidation
- Virtualization acceleration
- Power efficient
- Scales to tens-of-thousands of nodes

KEY FEATURES

- Virtual Protocol Interconnect
- Up to FDR 56Gb/s InfiniBand or 40/56 Gigabit Ethernet per port
- Single- and Dual-Port options available
- PCI Express 3.0 (up to 8GT/s)
- OCP Specification 2.0
- 1us MPI ping latency
- Hardware Offloads for NVGRE and VXLAN encapsulated traffic
- Application offload
- Data Center Bridging support
- GPU communication acceleration
- Precision Clock Synchronization
- Traffic steering across multiple cores
- Hardware-based I/O virtualization
- End-to-end QoS and congestion control
- Intelligent interrupt coalescence
- Advanced Quality of Service
- RoHS-R6

over Ethernet networks. Leveraging Data Center Bridging capabilities, RoCE provides efficient low latency RDMA services over Layer 2 and Layer 3 Ethernet. With link-level interoperability in existing Ethernet infrastructure, Network Administrators can leverage existing data center fabric management solutions.

Virtualized Overlay Networks – Infrastructure as a Service (IaaS) cloud demands that data centers host and serve multiple tenants, each with its own isolated network domain over a shared network infrastructure. To achieve maximum efficiency, data center operators create overlay networks that carry traffic from individual Virtual Machines (VMs) in encapsulated formats such as NVGRE and VXLAN over a logical "tunnel," thereby stretching a virtual layer-2 network over the physical layer-3 network. Overlay Network architecture introduces an additional layer of packet processing at the hypervisor level, adding and removing protocol headers for the encapsulated traffic. The new encapsulation prevents many of the traditional "offloading" capabilities (for example, checksum and LSO) from being performed at the adapter card.

ConnectX-3 Pro effectively addresses the increasing demand for an overlay network, enabling superior performance by introducing advanced NVGRE and VXLAN hardware offload engines that allow traditional offloads to be performed on the encapsulated traffic. With ConnectX-3 Pro, data center operators can decouple the overlay network layer from the physical adapter card performance, thus achieving native performance in the new network architecture.

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 Pro reduce the CPU overhead of IP packet transport. Sockets acceleration software further increases performance for latency sensitive applications.

I/O Virtualization – ConnectX-3 Pro SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines within the server. I/O virtualization with ConnectX-3 Pro gives data center managers better server utilization while reducing cost, power, and cable complexity.

Quality of Service – Resource allocation per application or per VM is provided and protected by the advanced QoS supported by ConnectX-3 Pro EN. Service levels for multiple traffic types can be based on IETF DiffServ or IEEE 802.1p/Q allowing system administrators to prioritize traffic by application, virtual machine, or protocol. This powerful combination of QoS and prioritization provides the ultimate fine-grained control of traffic – ensuring that applications run smoothly in today's complex environments.

Storage Accelerated – 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.

Software Support

All Mellanox adapter cards are supported by a Windows, Linux distributions, VMware, and Citrix XenServer. ConnectX-3 Pro 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 802.3ad 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
- Jumbo frame support (9600B)

HARDWARE-BASED I/O VIRTUALIZATION
- Single Root IOV (SR-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 (RoCE)
- TCP/UDP/IP stateless offload
- Intelligent interrupt coalescence

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

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

COMPLIANCE

GENERAL
- Adapters for Open Compute Project (OCP) Specification 2.0

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.
- Passive copper cable with ESD protection
- Powered connectors for optical and active cable support
- QSFP to SFP+ connectivity through QSA module

OPERATING SYSTEMS/DISTRIBUTIONS
- Citrix XenServer 6.1
- Novell SLES, Red Hat Enterprise Linux (RHEL), Ubuntu and other Linux distributions
- Microsoft Windows Server 2008/2012/2012 R2
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF)
- VMware ESXi

Ordering Part Number	Ethernet Ports	Dimensions
MCX345A-FCPN	ConnectX®-3 Pro VPI network interface card for OCP, FDR and 40/56GbE single-port QSFP, PCIe3.0 x8, no bracket, RoHS R6	11cm x 6.8cm
MCX346A-FCPN	ConnectX®-3 Pro VPI network interface card for OCP, FDR and 40/56GbE dual-port QSFP, PCIe3.0 x8, no bracket, RoHS R6	11cm x 6.8cm

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

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

MLNX-15-4719 Rev 1.0