

ConnectX[®]-3 Pro EN for Open Compute Project Supporting OCP Specification 2.0

Single/Dual-Port 40 Gigabit Ethernet Adapters with PCI Express 3.0

Mellanox ConnectX-3 Pro EN 40 Gigabit Ethernet Network Interface Card (NIC) with PCI Express 3.0 delivers high-bandwidth, low latency and industry-leading Ethernet connectivity for Open Compute Project (OCP) server and storage applications in Web 2.0, Enterprise Data Centers and Cloud infrastructure.

Web2.0, public and private clouds, clustered databases, web infrastructure, and high frequency trading are just a few applications that will achieve significant throughput and latency improvements resulting in faster access, real-time response and more virtual machines hosted per server. ConnectX-3 Pro 40GbE for Open Compute Project (OCP) specification 2.0 improves network performance by increasing available bandwidth while decreasing the associated transport load on the CPU especially in virtualized server environments.

World-Class Ethernet Performance

Virtualized Overlay Networks –

Infrastructure as a Service (IaaS) cloud demands that data centers host and serve multiple tenants, each with their 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 (e.g. checksum, LSO) from being performed at the NIC.


ConnectX-3 Pro EN 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 EN, data center operators can decouple the overlay network layer from the physical NIC performance, thus achieving native performance in the new network architecture.

RDMA over Converged Ethernet –

ConnectX-3 Pro EN utilizing IBTA RoCE technology provides efficient RDMA (Remote Direct Memory Access) services, delivering low-latency and high-performance to bandwidth and latency sensitive applications. 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 40GbE. The hardware-based stateless offload and flow steering engines in ConnectX-3 Pro EN reduce the CPU overhead of IP packet transport, freeing more processor cycles to work on the application. Sockets acceleration software further increases performance for latency sensitive applications.

HIGHLIGHTS

BENEFITS

- 40Gb/s connectivity for servers and storage
- Open Compute Project Form Factor
- Industry-leading throughput and latency performance
- Cutting edge performance in virtualized overlay networks (VXLAN and NVGRE)
- I/O consolidation
- Virtualization acceleration
- Software compatible with standard TCP/UDP/IP and iSCSI stacks

KEY FEATURES

- Single and Dual 40 Gigabit Ethernet ports
- PCI Express 3.0 (up to 8GT/s)
- OCP Specification 2.0
- Low Latency RDMA over Ethernet
- Hardware Offloads for NVGRE and VXLAN encapsulated traffic
- Data Center Bridging support
- TCP/IP stateless offload in hardware
- Traffic steering across multiple cores
- Hardware-based I/O
- End-to-end QoS and congestion control
- Virtualization
- Intelligent interrupt coalescence
- Advanced Quality of Service
- RoHS-R6

I/O Virtualization –

ConnectX-3 Pro EN with SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VMs) within the server. ConnectX-3 Pro EN gives data center managers better server utilization and LAN and SAN unification while reducing costs, power, and 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.

Software Support

ConnectX-3 Pro EN is supported by a full suite of software drivers for Microsoft Windows, Linux distributions, VMware and Citrix XENServer. ConnectX-3 Pro EN supports stateless offload and is fully interoperable with standard TCP/UDP/IP stacks. ConnectX-3 Pro EN supports various management interfaces and has a rich set of configuring and management tools across operating systems.

SPECIFICATIONS*

ETHERNET

- IEEE 802.3ae 10 Gigabit Ethernet
- IEEE Std 802.3ba 40 Gigabit Ethernet
- IEEE 802.3ad Link Aggregation and Failover
- IEEE 802.3az Energy Efficient Ethernet
- IEEE 802.1Q, .1p VLAN tags and priority
- IEEE 802.1Qau Congestion Notification
- IEEE Std 802.1Qbg
- IEEE P802.1Qaz D0.2 ETS
- IEEE P802.1Qbb D1.0 Priority-based Flow Control
- IEEE 1588v2
- Jumbo frame support (9KB)

OVERLAY NETWORKS

- VXLAN and NVGRE - A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks. Network Virtualization hardware offload engines

MANAGEMENT AND CONTROL INTERFACES

- NC-SI, MCTP over SMBus - Baseboard Management Controller interface

HARDWARE-BASED I/O VIRTUALIZATION

- Single Root IOV
- Address translation and protection
- Dedicated adapter resources
- Enhanced QoS for vNICs
- VMware NetQueue support

ADDITIONAL CPU OFFLOADS

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

FLEXBOOT™ TECHNOLOGY

- Remote boot over Ethernet
- Remote boot over iSCSI

PROTOCOL SUPPORT

- Open MPI, OSU MVAPICH, Intel MPI, MS
- MPI, Platform MPI
- TCP/UDP
- 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

CONNECTIVITY

- Interoperable with 10/40GigE switches
- QSFP+ connectors
- QSFP to SFP+ connectivity through QSA module
- Passive copper cable with ESD protection
- Powered connectors for optical and active cable support

OPERATING SYSTEMS/DISTRIBUTIONS

- Novell SLES, Red Hat Enterprise Linux (RHEL), Fedora, CentOS and other Linux distributions.
- Microsoft Windows Server
- VMware ESX Server
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF)

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

Ordering Part Number	Ethernet Ports	Dimensions
MCX345A-BCPN	40GbE single port QSFP+	11cm x 6.8cm
MCX345A-BCQN	40GbE single port QSFP+, with NC-SI host management protocol enabled	11cm x 6.8cm
MCX346A-BCPN	40GbE dual port QSFP+	11cm x 6.8cm
MCX346A-BCQN	40GbE dual port QSFP+, with NC-SI host management protocol enabled	11cm x 6.8cm



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