



## ADAPTER CARDS

# ConnectX<sup>®</sup>-2 EN 40G

## 40 Gigabit Ethernet Adapter with PCI Express 2.0

ConnectX-2 EN 40 Gigabit Ethernet Network Interface Card (NIC) delivers breakthrough connectivity for servers and storage. IT manager will be able to build out high-bandwidth fabrics for next generation data centers enabling the highest platform efficiency and CPU utilization. Virtualized environments will benefit as well from the hardware-based accelerations enabling more virtual machines per server and superior application performance.

The single port ConnectX-2 EN 40G adapter comes with a QSFP connector suitable for use with copper or fiber optic cables.

### World-Class Ethernet Performance

ConnectX-2 EN 40G with PCI Express 2.0 delivers industry-leading Ethernet bandwidth required by multi-core, multi-socket platforms. Hardware-based stateless offload engines handle the TCP/UDP/IP segmentation, reassembly, and checksum calculations that would otherwise burden the host. Total cost of ownership is optimized by maintaining an end-to-end Ethernet network on existing operating systems and applications.

### Converged Ethernet

ConnectX-2 EN 40G delivers the features needed for a converged network with support for Data Center Bridging (DCB). Fibre Channel frame encapsulation compliant with T11 and hardware offloads simplifies FCoE deployment. Low Latency Ethernet (LLE) running over DCB fabrics provides optimal performance for RDMA and Send/Receive services. By maintaining link-level interoperability with existing Ethernet networks, IT managers can leverage existing data center fabric management solutions.

### I/O Virtualization


ConnectX-2 EN 40G supports hardware-based I/O virtualization, providing dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server. ConnectX-2 EN 40G 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 by the advanced QoS supported by ConnectX-2 EN 40G. Service levels for multiple traffic types can be based on IETF DiffServ or IEEE 802.1p/Q, along with the Data Center Bridging enhancements, allowing system administrators to prioritize traffic by application, virtual machine, or protocol. This powerful combination of QoS and prioritization provides the ultimate fine-grain control of traffic – ensuring that applications run smoothly in today's complex environment.

### Software Support

ConnectX-2 EN 40G is supported by a full suite of software drivers for Microsoft Windows, Linux distributions, VMware and Citrix XEN Server. ConnectX-2 EN 40G supports stateless offload and is fully interoperable with standard TCP/UDP/IP stacks. ConnectX-2 EN 40G supports various



### BENEFITS

- 40Gb/s connectivity for servers and storage
- Industry-leading throughput and latency performance
- I/O consolidation
- Virtualization acceleration
- High-performance networking and storage access
- Software compatible with standard TCP/UDP/IP and iSCSI stacks

### KEY FEATURES

- Single 10GBASE-CR4 port
- PCI Express (up to 5GT/s)
- Low Latency Ethernet
- Data Center Bridging support
- T11 FCoE frame support
- TCP/UDP/IP stateless offload in hardware
- Traffic steering across multiple cores
- Hardware-based I/O virtualization
- Intelligent interrupt coalescence
- Advanced Quality of Service
- RoHS-R6

management interfaces and has a rich set of configuring and management tools across operating systems.

## FEATURE SUMMARY

### ETHERNET

- IEEE Draft P802.3ba/D2.0 40GBASE-CR4
- IEEE Std 802.3x Pause
- IEEE Std 802.1Q VLAN tags, .1p Priorities
- IEEE P802.1au D2.0 Congestion Notification
- IEEE P802.1az D0.2 Enhanced Transmission Selection
- IEEE P802.1bb D1.0 Priority-based Flow Control
- Multicast
- Jumbo frame support (10KB)
- 128 MAC/VLAN addresses per port

### TCP/UDP/IP STATELESS OFFLOAD

- TCP/UDP/IP checksum offload
- TCP Large Send (< 64KB) or Giant Send (64KB-16MB) Offload for segmentation
- Receive Side Scaling (RSS) up to 32 queues
- Line rate packet filtering

### ADDITIONAL CPU OFFLOADS

- RDMA, Send/Receive (LLE)
- Traffic steering across multiple cores
- Intelligent interrupt coalescence
- Compliant to Microsoft RSS and NetDMA

### HARDWARE-BASED I/O VIRTUALIZATION

- Single Root IOV
- Address translation and protection
- Dedicated adapter resources and guaranteed isolation
- Multiple queues per virtual machine
- Hardware switching between guest OSs
- Enhanced QoS for vNICs
- VMware NetQueue support

### STORAGE SUPPORTS

- Fibre Channel over Ethernet
- T11-compliant frame format

## 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
- 5.0GT/s link rate x8 (40+40Gb/s bidirectional bandwidth)
- Fits x8 or x16 slots
- Support for MSI/MSI-X mechanisms

### CONNECTIVITY

- 7m+ of passive copper cable
- Active cable support

### OPERATING SYSTEMS/DISTRIBUTIONS

- Novell SuSE Linux Enterprise Server (SLES), Red Hat Enterprise Linux (RHEL), and other Linux distributions
- Microsoft Windows XP, Server 2003 / 2008
- VMware ESX 3.5
- Citrix XenServer 4.1

### MANAGEMENT

- MIB, MIB-II, MIB-II Extensions, RMON, RMON 2
- Configuration and diagnostic tools

## Adapter Cards

Ordering Part Number	Ethernet Ports	Host Bus	Power (Typ)	Dimensions w/o Brackets
MNQH19-XTC	Single QSFP with powered connector	PCIe 2.0 5.0GT/s	TBD	13.6cm x 6.4cm



350 Oakmead Pkwy, Suite 100, Sunnyvale, CA 94085  
 Tel: 408-970-3400 • Fax: 408-970-3403  
[www.mellanox.com](http://www.mellanox.com)

© Copyright 2009, Mellanox Technologies. All rights reserved.  
 Mellanox, ConnectX, InfiniBlast, InfiniBridge, InfiniHost, InfiniRISC, InfiniScale, and InfiniPCI are registered trademarks of Mellanox Technologies, Ltd. BridgeX, PhyX, and Virtual Protocol Interconnect are trademarks of Mellanox Technologies, Ltd. All other trademarks are property of their respective owners.