




ConnectX[®]-4 Lx EN Card

Up to 50Gb/s Ethernet Adapter Cards

1/10/25/40/50 Gigabit Ethernet adapter cards supporting RDMA, Overlay Network Encapsulation/Decapsulation and more

ConnectX-4 Lx EN adapter cards with 1/10/25/40/50 Gb/s Ethernet connectivity addresses virtualized infrastructure challenges for today's demanding markets and applications. Providing true hardware-based I/O isolation with unmatched scalability and efficiency, ConnectX-4 Lx EN provides a cost-effective and flexible Ethernet adapter solution for Web 2.0, cloud, data analytics, database, and storage platforms.

With the exponential increase in usage of data and the creation of new applications, the demand for the highest throughput, lowest latency, virtualization and sophisticated data acceleration engines continues to rise. ConnectX-4 Lx EN adapter cards enable data centers to leverage leading interconnect adapters for increasing their operational efficiency, improving servers' utilization, maximizing applications productivity, while reducing total cost of ownership (TCO).

ConnectX-4 Lx EN adapter cards provide a combination of 1, 10, 25, 40, and 50 GbE bandwidth, sub-microsecond latency and a 75 million packets per second message rate. They include native hardware support for RDMA over Converged Ethernet (RoCE), Ethernet stateless offload engines, Overlay Networks, GPUDirect[®] technology and Mellanox Multi-Host[®] technology.

Mellanox Multi-Host[®] Technology

Innovative Mellanox Multi-Host technology enables data centers to design and build scale-out heterogeneous compute and storage racks, with direct connectivity between compute elements and the network. Significantly improving cost savings, flexibility, and total cost of ownership, Mellanox Multi-Host technology provides better power and performance, while achieving maximum data processing and data transfer at minimum capital and operational expenses.

Mellanox Multi-Host works by allowing multiple hosts to connect into a single interconnect adapter, by separating the adapter PCIe interface into several independent PCIe interfaces. Each interface connects to a separate host CPU—with no performance degradation. Reducing data center CAPEX and OPEX, Mellanox Multi-Host slashes switch port management and power usage by reducing the number of cables, NICs and switch ports required by four independent servers, from four to one of each. Additional features & benefits of Mellanox Multi-Host technology:

- Enables IT managers to remotely control the configuration and power state of each host individually; guaranteeing host security and isolation, the management of one host does not affect host traffic performance nor the management of other hosts.
- Lowering total cost of ownership (TCO), Mellanox Multi-Host uses a single BMC, with independent NC-SI/MCTP management channels for each of the managed hosts.
- Mellanox Multi-Host also supports a heterogeneous data center architecture; the various hosts connected to the single adapter can be x86, Power, GPU, Arm or FPGA, thereby removing any limitations in passing data or communicating between compute elements.

ConnectX[®]-4 Lx

HIGHLIGHTS

FEATURES

- 1/10/25/40/50 Gb/s speeds
- Single and dual-port options
- Virtualization
- Low latency RDMA over Converged Ethernet (RoCE)
- Mellanox Multi-Host technology connects up to 4 independent hosts
- CPU offloading of transport operations
- Application offloading
- Mellanox PeerDirect[®] communication acceleration
- Hardware offloads for NVGRE, VXLAN and GENEVE encapsulated traffic
- End-to-end QoS and congestion control
- Hardware-based I/O virtualization
- RoHS compliant
- ODCC compatible
- Various form factors available

BENEFITS

- High performance boards for applications requiring high bandwidth, low latency and high message rate
- Industry leading throughput and latency for Web 2.0, Cloud and Big Data applications
- Smart interconnect for x86, Power, ARM, and GPU-based compute and storage platforms
- Cutting-edge performance in virtualized overlay networks
- Efficient I/O consolidation, lowering data center costs and complexity
- Virtualization acceleration
- Power efficiency

Wide Selection of Ethernet Adapter Cards

ConnectX-4 Lx EN adapter cards offer a cost-effective Ethernet adapter solution for 1, 10, 25, 40 and 50 Gb/s Ethernet speeds, enabling seamless networking, clustering, or storage. The adapter reduces application runtime, and offers the flexibility and scalability to make infrastructure run as efficiently and productively as possible.

ConnectX-4 Lx Ethernet adapter cards are available in several form factors including: low-profile stand-up PCIe, OCP 2.0 Type 1, OCP 2.0 Type 2, and OCP 3.0 small form factor. (See the portfolio on the last page.)

I/O Virtualization

ConnectX-4 Lx EN SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VMs) within the server. I/O virtualization with ConnectX-4 Lx EN gives data center administrators better server utilization while reducing cost, power, and cable complexity, allowing more virtual machines and more tenants on the same hardware.

Overlay Networks

In order to better scale their networks, data center operators often create overlay networks that carry traffic from individual virtual machines over logical tunnels in encapsulated formats such as NVGRE and VXLAN. While this solves network scalability issues, it hides the TCP packet from the hardware offloading engines, placing higher loads on the host CPU. ConnectX-4 Lx EN effectively addresses this by providing advanced NVGRE, VXLAN and GENEVE hardware offloading engines that encapsulate and de-capsulate the overlay protocol headers, enabling the traditional offloads to be performed on the encapsulated traffic for these and other tunneling protocols (GENEVE, MPLS, QinQ, and so on). With ConnectX-4 Lx EN, data center operators can achieve native performance in the new network architecture.

RDMA over Converged Ethernet (RoCE)

ConnectX-4 Lx EN supports RoCE specifications delivering low-latency and high-performance over Ethernet networks. Leveraging data center bridging (DCB) capabilities as well as ConnectX-4 Lx EN advanced congestion control hardware mechanisms, RoCE provides efficient low-latency RDMA services over Layer 2 and Layer 3 networks.

Mellanox PeerDirect®

Mellanox PeerDirect communication provides high efficiency RDMA access by eliminating unnecessary internal data copies between components on the PCIe bus (for example, from GPU to CPU), and therefore significantly reduces application run time. ConnectX-4 Lx EN advanced acceleration technology enables higher cluster efficiency and scalability to tens of thousands of nodes.

Storage Acceleration

Storage applications will see improved performance with the higher bandwidth ConnectX-4 Lx EN delivers. Moreover, standard block and file access protocols can leverage RoCE for high-performance storage access. A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks.

Host Management

Mellanox host management and control capabilities include NC-SI over MCTP over SMBus, and MCTP over PCIe - Baseboard Management Controller (BMC) interface, as well as PLDM for Monitoring and Control DSP0248 and PLDM for Firmware Update DSP0267.

Software Support

All Mellanox adapter cards are supported by Windows, Linux distributions, VMware, FreeBSD, and Citrix XENServer. ConnectX-4 Lx EN supports various management interfaces and has a rich set of tools for configuration and management across operating systems.

Additionally, ConnectX-4 Lx EN provides the option for a secure firmware update check using digital signatures to prevent remote attackers from uploading malicious firmware images; this ensures that only officially authentic images produced by Mellanox can be installed, regardless whether the source of the installation is the host, the network, or a BMC.

Features*

Ethernet

- 50GbE / 40GbE / 25GbE / 10GbE / 1GbE
- IEEE 802.3bj, 802.3bm 100 Gigabit Ethernet
- 25G Ethernet Consortium 25, 50 Gigabit Ethernet
- IEEE 802.3ba 40 Gigabit Ethernet
- IEEE 802.3ae 10 Gigabit Ethernet
- IEEE 802.3az Energy Efficient Ethernet
- IEEE 802.3ap based auto-negotiation and KR startup
- Proprietary Ethernet protocols (20/40GBASE-R2, 50GBASE-R4)
- IEEE 802.3ad, 802.1AX Link Aggregation
- IEEE 802.1Q, 802.1P VLAN tags and priority
- IEEE 802.1Qau (QCN) – Congestion Notification
- IEEE 802.1Qaz (ETS)
- IEEE 802.1Qbb (PFC)
- IEEE 802.1Qbg
- IEEE 1588v2
- Jumbo frame support (9.6KB)

Mellanox Multi-Host

- Up to 4 separate PCIe interfaces to 4 independent hosts
- Two PCIe x4 to two hosts, or four PCIe x4 to four hosts, or four PCIe x2 to four hosts

- Independent NC-SI SMBus interfaces
- Independent stand-by and wake-on-LAN signals

Enhanced Features

- Hardware-based reliable transport
- Collective operations offloads
- Vector collective operations offloads
- PeerDirect RDMA (aka GPUDirect® communication acceleration)
- 64/66 encoding
- Extended Reliable Connected transport (XRC)
- Dynamically Connected transport (DCT)
- Enhanced Atomic operations
- Advanced memory mapping support, allowing user mode registration and remapping of memory (UMR)
- On demand paging (ODP) – registration free RDMA memory access

CPU Offloads

- RDMA over Converged Ethernet (RoCE)
- TCP/UDP/IP stateless offload
- LSO, LRO, checksum offload
- RSS (can be done on encapsulated packet), TSS, VLAN insertion/stripping, Receive flow steering
- Intelligent interrupt coalescence

Overlay Networks

- Stateless offloads for overlay networks and tunneling protocols
- Hardware offload of encapsulation and decapsulation of NVGRE and VXLAN overlay networks

Hardware-Based I/O Virtualization

- Single Root IOV
- Multi-function per port
- Address translation and protection
- Multiple queues per virtual machine
- Enhanced QoS for vNICs
- VMware NetQueue support

Virtualization

- SR-IOV: Up to 256 Virtual Functions
- SR-IOV: Up to 8 Physical Functions per port
- Virtualization hierarchies (e.g. NPAR)
 - Virtualizing Physical Functions on a physical port
 - SR-IOV on every Physical Function
- Ingress and egress QoS levels
- Guaranteed QoS for VMs

Protocol Support

- OpenMPI, IBM PE, OSU MPI (MVAPICH/2), Intel MPI
- Platform MPI, UPC, Open SHMEM
- TCP/UDP, MPLS, VxLAN, NVGRE, GENEVE
- iSER, NFS RDMA, SMB Direct
- uDAPL

Management and Control Interfaces

- NC-SI over MCTP over SMBus and NC-SI over MCTP over PCIe - Baseboard Management Controller interface
- SDN management interface for managing the eSwitch
- I²C interface for device control and configuration
- General Purpose I/O pins
- SPI interface to flash
- JTAG IEEE 1149.1 and IEEE 1149.6

Remote Boot

- Remote boot over Ethernet
- Remote boot over iSCSI
- PXE and UEFI

* This section describes hardware features and capabilities. Please refer to the driver and firmware release notes for feature availability.

Compatibility

PCI Express Interface

- PCIe Gen 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

Operating Systems/Distributions*

- RHEL/CentOS
- Windows
- FreeBSD
- VMware
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF-2)

Connectivity

- Interoperable with 1/10/25/40/50 Gb/s Ethernet switches
- Passive copper cable with ESD protection
- Powered connectors for optical and active cable support

Standards*

- IEEE 802.3cd, 50,100 and 200 Gb/s
- IEEE 802.3by, 25, 50 Gb/s supporting all FEC modes
- IEEE 802.3ba 40 Gigabit Ethernet
- IEEE 802.3ae 10 Gigabit Ethernet
- IEEE 802.3az Energy Efficient Ethernet (supports only "Fast-Wake" mode)
- IEEE 802.3ap based auto-negotiation and KR startup
- IEEE 802.3ad, 802.1AX Link Aggregation
- IEEE 802.1Q, 802.1P VLAN tags and priority
- IEEE 802.1Qaz (ETS)
- IEEE 802.1Qbb (PFC)
- IEEE 802.1Qbg
- IEEE 1588v2
- 25G/50G Ethernet Consortium "Low Latency FEC" for 50/100/200GE PAM4 links
- PCI Express Gen 3.0

Card Portfolio & Ordering Information

Table 1 - PCIe HHL Form Factor

Max Network Speed	Interface Type	Supported Ethernet Speeds (GbE)	Host Interface	Additional Features	OPN
1x 10GbE	SFP28	10, 1	PCIe 3.0 x8		MCX4111A-XCAT
2x 10GbE	SFP28	10, 1	PCIe 3.0 x8		MCX4121A-XCAT
			PCIe 3.0 x8	Host Management, UEFI Enabled	MCX4121A-XCHT
1x 25GbE	SFP28	25, 10, 1	PCIe 3.0 x8		MCX4111A-ACAT
			PCIe 3.0 x8	UEFI Enabled	MCX4111A-ACUT
2x 25GbE	SFP28	25, 10, 1	PCIe 3.0 x8		MCX4121A-ACAT
			PCIe 3.0 x8	UEFI Enabled	MCX4121A-ACUT
1x 40GbE	QSFP28	40, 25, 10, 1	PCIe 3.0 x8		MCX4131A-BCAT
1x 50GbE	QSFP28	50, 40, 25, 10, 1	PCIe 3.0 x8		MCX4131A-GCAT

Notes: All tall-bracket adapters are shipped with the tall bracket mounted and a short bracket as an accessory.
Dimensions without bracket: 14.2cm x 6.9cm (low profile)




Table 2 - OCP 2.0 Type 1 Form Factor

Max Network Speed	Interface Type	Supported Ethernet Speeds (GbE)	Host Interface	Additional Features	OPN
2x 10GbE	SFP28	10, 1	PCIe 3.0 x8	Host Management	MCX4421A-XCQN
1x 25GbE	SFP28	25, 10, 1	PCIe 3.0 x8		MCX4411A-ACAN
			PCIe 3.0 x8	Host Management	MCX4411A-ACQN
			PCIe 3.0 x8	UEFI Enabled	MCX4411A-ACUN
2x 25GbE	SFP28	25, 10, 1	PCIe 3.0 x8		MCX4421A-ACAN
			PCIe 3.0 x8	UEFI Enabled	MCX4421A-ACUN
			PCIe 3.0 x8	Host Management	MCX4421A-ACQN
1x 50GbE	QSFP28	50, 40, 25, 10, 1	PCIe 3.0 x8	Host Management	MCX4431A-GCAN
			PCIe 3.0 x8	Host Management, UEFI Enabled	MCX4431A-GCUN





Table 3 - OCP 2.0 Type 2 Form Factor

Max Network Speed	Interface Type	Supported Ethernet Speeds (GbE)	Host Interface	Additional Features	OPN
1x 40GbE	QSFP28	40, 25, 10, 1	PCIe 3.0 x8	Mellanox Multi-Host, Host Management	MCX4431M-BCAN
1x 50GbE	QSFP28	50, 40, 25, 10, 1	PCIe 3.0 x8	Mellanox Multi-Host, Host Management	MCX4431M-GCAN

Note: OCP2.0 cards are shipped without a bracket.

Table 4 - OCP 3.0 Small Form Factor

Max Network Speed	Interface Type	Supported Ethernet Speeds (GbE)	Host Interface	Additional Features	OPN
2x 25GbE	SFP28	25, 10, 1	PCIe 3.0 x8	Host Management, Thumbscrew bracket	MCX4621A-ACAB



† For illustration only. Actual products may vary.