ConnectX®–4 EN

Single/Dual-Port 100 Gigabit Ethernet Adapter Cards

ConnectX-4 EN Network Controller with 100Gb/s Ethernet connectivity, provide the highest performance and most flexible solution for high performance, Web 2.0, Cloud, data analytics, database, and storage platforms.

With the exponential growth of data being shared and stored by applications and social networks, the need for high-speed and high performance compute and storage data centers is skyrocketing.

ConnectX®-4 EN provides exceptional high performance for the most demanding data centers, public and private clouds, Web2.0 and BigData applications, and Storage systems, enabling today's corporations to meet the demands of the data explosion.

ConnectX-4 EN provides an unmatched combination of 100Gb/s bandwidth in a single port, the lowest available latency, and specific hardware offloads, addressing both today's andthe next generation's compute and storage data center demands.

I/O Virtualization

ConnectX-4 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 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 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. With ConnectX-4, data center operators can achieve native performance in the new network architecture.

ASAP²™

Mellanox ConnectX-4 EN offers Accelerated Switching And Packet Processing (ASAP²) technology to perform offload activities in the hypervisor, including data path, packet parsing, VxLAN and NVGRE encapsulation/decapsulation, and more.

ASAP² allows offloading by handling the data plane in the NIC hardware using SR-IOV, while maintaining the control plane used in today's software-based solutions unmodified. As a result, there is significantly higher performance without the associated CPU load. ASAP² has two formats: ASAP² Flex™ and ASAP² Direct™.

One example of a virtual switch that ASAP² can offload is OpenVSwitch (OVS).

HIGHLIGHTS

BENEFITS

– Highest performing silicon for applications requiring high bandwidth, low latency and high message rate
– World-class cluster, network, and storage performance
– Smart interconnect for x86, Power, ARM, and GPU-based compute and storage platforms
– Cutting-edge performance in virtualized overlay networks VXLAN, NVGRE, GENEVE
– Efficient I/O consolidation, lowering data center costs and complexity
– Virtualization acceleration
– Power efficiency
– Scalability to tens-of-thousands of nodes

KEY FEATURES

– 100Gb/s Ethernet per port
– 1/10/20/25/40/50/56/100Gb/s speeds
– Single and dual-port options available
– Erasure Coding offload
– T10-DIF Signature Handover
– Accelerated Switching and Packet Processing (ASAP²)
– 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-R6

† For illustration only. Actual products may vary.

RDMA over Converged Ethernet (RoCE)

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

Mellanox PeerDirect™

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 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 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.

Erasure Coding Offload

ConnectX-4 EN delivers advanced Erasure Coding offloading capability, enabling distributed Redundant Array (RAID) of Inexpensive Disks, a data storage technology that combines multiple disk drive components into a logical unit for the purposes of data redundancy and performance improvement. ConnectX-4 EN's Reed-Solomon capability introduces redundant block calculations, which, together with RDMA, achieves high performance and reliable storage access.

Signature Handover

ConnectX-4 EN supports hardware checking of T10 Data Integrity Field / Protection Information (T10-DIF/PI), reducing the CPU overhead and accelerating delivery of data to the application. Signature handover is handled by the adapter on ingress and/or egress packets, reducing the load on the CPU at the Initiator and/or Target machines.

Software Support

All Mellanox adapter cards are supported by Windows, Linux distributions, VMware, FreeBSD, and Citrix XENServer. ConnectX-4 EN adapters support OpenFabrics-based RDMA protocols and software and are compatible with configuration and management tools from OEMs and operating system vendors.

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 x16
– Auto-negotiates to x16, x8, x4, x2, or x1
– Support for MSI/MSI-X mechanisms

CONNECTIVITY
– Interoperable with 1/10/25/40/50/100Gb Ethernet switches
– Passive copper cable with ESD protection
– Powered connectors for optical and active cable support

OPERATING SYSTEMS/DISTRIBUTIONS*
– RHEL/CentOS
– Windows
– FreeBSD
– VMware
– OpenFabrics Enterprise Distribution (OFED)
– OpenFabrics Windows Distribution (WinOF)

Ordering Part Number	Description	Dimensions w/o Brackets
MCX415A-CCAT	ConnectX-4 EN network interface card, 100GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)
MCX416A-CCAT	ConnectX-4 EN network interface card, 100GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)
MCX415A-BCAT	ConnectX-4 EN network interface card, 40/56GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)
MCX416A-BCAT	ConnectX-4 EN network interface card, 40/56GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)
MCX415A-GCAT	ConnectX-4 EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)
MCX416A-GCAT	ConnectX-4 EN network interface card, 50GbE dual-port QSFP28, PCIe3.0 x16, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)
MCX413A-BCAT	ConnectX-4 EN network interface card, 40/56GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)
MCX414A-BCAT	ConnectX-4 EN network interface card, 40/56GbE dual-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)
MCX413A-GCAT	ConnectX-4 EN network interface card, 50GbE single-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)
MCX414A-GCAT	ConnectX-4 EN network interface card, 50GbE dual-port QSFP28, PCIe3.0 x8, tall bracket, ROHS R6	14.2cm x 6.9cm (Low Profile)

FEATURE SUMMARY*

ETHERNET
- 100GbE / 56GbE / 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, 50/56GBASE-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)

ENHANCED FEATURES
- Hardware-based reliable transport
- Collective operations offloads
- Vector collective operations offloads
- Mellanox 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

STORAGE OFFLOADS
- RAID offload - erasure coding (Reed-Solomon) offload
- T10 DIF - Signature handover operation at wire speed, for ingress and egress traffic

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 16 Physical Functions per port
- Virtualization hierarchies (e.g. NPAR)
  » Virtualizing Physical Functions on a physical port
  » SR-IOV on every Physical Function
- 1K ingress and egress QoS levels
- Guaranteed QoS for VMs

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

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

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, MCTP over SMBus and 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

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

MLNX-15-4392 Rev 1.6