



ConnectX[®]-6 HDR100 Single Port Adapter Card

HDR100 InfiniBand Adapter Card



The HDR100 InfiniBand Adapter Card offers industry-leading performance, smart offloads and In-Network Computing, leading to the highest return on investment for High-Performance Computing, Cloud, Web 2.0, Storage and Machine Learning applications. It is ideally suited to support the Dell PowerEdge R640 and PowerEdge C6420 servers and extends the Dell EMC InfiniBand adapter portfolio with enhanced performance and improved security. ConnectX-6 HDR-100 adapters support bandwidth of 100Gb/s HDR InfiniBand speeds. One of the benefits of HDR100 is the use of breakout cables on HDR 200Gb ports, reducing the number of switches a customer needs in their HPC environment and reducing TCO.

HPC ENVIRONMENTS

Over the past decade, Mellanox has consistently driven HPC performance to new record heights. With the introduction of the ConnectX-6 HDR100 single port Adapter card, Mellanox continues to pave the way with new features and unprecedented performance for the HPC market.

ConnectX-6 supports the evolving Co-Design paradigm, which transforms the network into a distributed processor. With its In-Network Computing and In-Network Memory capabilities, ConnectX-6 offloads computation even further to the network, saving CPU cycles and increasing network efficiency.

ConnectX-6 HDR100 InfiniBand single port Adapter Card utilizes IBTA RDMA (Remote Data Memory Access) technology, delivering low-latency and high performance. ConnectX-6 enhances RDMA network capabilities even further by delivering end-to-end packet level flow control.

MACHINE LEARNING AND BIG DATA ENVIRONMENTS

Data analytics has become an essential function within many enterprise data centers, clouds and Hyperscale platforms. Machine learning relies on especially high throughput and low latency to train deep neural networks and to improve recognition and classification accuracy. ConnectX-6 is the perfect solution to provide machine learning applications with the levels of performance and scalability that they require.

ConnectX-6 utilizes the RDMA technology to deliver low-latency and high performance. ConnectX-6 enhances RDMA network capabilities even further by delivering end-to-end packet level flow control.

CLOUD AND WEB 2.0 ENVIRONMENTS

Telco, Cloud and Web 2.0 customers developing their platforms on Software Defined Network (SDN) environments are leveraging the Virtual Switching capabilities of the Operating Systems on their servers to enable maximum flexibility in the management and routing protocols of their networks.

The vSwitch/vRouter offload functions supported by ConnectX-5 and ConnectX-6 include encapsulation and de-capsulation of overlay network headers, as well as stateless offloads of inner packets, packet headers re-write (enabling NAT functionality), hairpin, and more.

In addition, ConnectX-6 offers intelligent flexible pipeline capabilities, including programmable flexible parser and flexible match-action tables, which enable hardware offloads for future protocols.

HIGHLIGHTS

FEATURES

- Max bandwidth of 100Gb/s
- Up to 215 million messages/sec
- Sub 0.6usec latency
- Block-level XTS-AES mode hardware encryption
- FIPS capable
- Advanced storage capabilities including block-level encryption and checksum offloads
- Supports both 50G SerDes (PAM4) and 25 SerDes (NRZ)-based ports
- Best-in-class packing with nsub-nanosecond accuracy
- PCIe Gen3 support
- RoHS compliant
- ODCC compatible

BENEFITS

- Industry-leading throughput, low CPU utilization and high message rate
- Highest performance and most intelligent fabric for compute and storage infrastructures
- Cutting-edge performance in virtualized networks including Network Function Virtualization (NFV)
- Smart interconnect for x86, Arm, GPU and FPGA-based compute and storage platforms
- Flexible programmable pipeline for new network flows
- Cutting-edge performance in virtualized networks, e.g., NFV
- Efficient service chaining enablement
- Increased I/O consolidation efficiencies, reducing data center costs & complexity

STANDARD HOST MANAGEMENT

Mellanox host management and control capabilities include NC-SI over MCTP over SMBus, and MCTP over PCIe - Baseboard Management Controller (BMC) interface.

SOCKET DIRECT

Mellanox's Socket Direct technology improves the performance of dual-socket servers in numerous ways, such as by enabling each of their CPUs to access the

network through a dedicated PCIe interface. As the connection from each CPU to the network bypasses the QPI (UPI) and the second CPU, Socket Direct reduces latency and CPU utilization. Moreover, each CPU handles only its own traffic (and not that of the second CPU), thus optimizing CPU utilization even further.

Socket Direct also enables GPUDirect® RDMA for all CPU/GPU pairs by ensuring that GPUs are linked to the CPUs closest to the adapter card. Socket Direct enables Intel® DDIO optimization on both sockets by creating a direct connection between the sockets and the adapter card.

COMPATIBILITY*

PCI Express Interface

- PCIe Gen 3.0, 2.0, 1.1 compatible
- 2.5, 5.0, 8GT/s link rate
- 16-lanes of PCIe
- Support for PCIe x1, x2, x4, x8, and x16 configurations
- PCIe Atomic
- TLP (Transaction Layer Packet) Processing Hints (TPH)

- PCIe switch Downstream Port Containment (DPC) enablement for PCIe hot-plug
- Advanced Error Reporting (AER)
- Access Control Service (ACS) for Peer-to-Peer secure communication
- Process Address Space ID (PASID) Address Translation Services (ATS)
- Support for MSI/MSI-X mechanisms

Operating Systems/Distributions*

- RHEL, SLES, and other major Linux distributions
- Windows
- FreeBSD
- VMware
- OpenFabrics Enterprise Distrib.(OFED)
- OpenFabrics Windows Distrib. (WinOF-2)

Connectivity

- Interoperability with InfiniBand switches (up to HDR100, as 2 lanes of 50Gb/s data rate)
- Passive copper cable with ESD protection
- Powered connectors for optical and active cable support

FEATURES*

InfiniBand

- HDR100 / EDR / FDR
- IBTA Specification 1.3 compliant
- RDMA, Send/Receive semantics
- Hardware-based congestion control
- Atomic operations
- 16 million I/O channels
- 256 to 4Kbyte MTU, 2Gbyte messages
- 8 virtual lanes + VL15

GPUDirect®) communication acceleration

- 64/66 encoding

CPU Offloads

- TCP/UDP/IP offloads
- Intelligent interrupt coalescence
- Header rewrite supporting hardware offload of NAT router

Storage Offloads

- NVMe over Fabric offloads for target machine
- Erasure Coding offload - offloading Reed-Solomon calculations
- T10 DIF - signature handover

operation at wire speed, for ingress and egress traffic

- Storage Protocols: SRP, iSER, NFS RDMA, SMB Direct

Hardware-Based I/O Virtualization

- Single Root IOV
- Address translation and protection
- SR-IOV: Up to 127 Virtual Functions
- Configurable and user-programmable QoS
- Guaranteed QoS for VMs

HPC Software Libraries

- HPC-X, OpenMPI, MVAPICH, MPICH,

OpenSHMEM, PGAS and varied commercial packages

Management and Control

- 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

Remote Boot

- Remote boot over InfiniBand
- Unified Extensible Firmware Interface (UEFI)
- Pre-execution Environment (PXE)

Enhanced Features

- Hardware-based reliable transport
- Collective operations offloads
- Vector collective operations offloads
- PeerDirect™ RDMA (aka

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

Table 1 - Part Numbers and Descriptions

New Mellanox Adapter	Tied SKU	Customer Kit SKU	DPN
Mellanox ConnectX-6 Single Port HDR100 QSFP56 InfiniBand Adapter, PCIe Full Height	540-BCMVM	540-BCMUM	Y1T43
Mellanox ConnectX-6 Single Port HDR100 QSFP56 InfiniBand Adapter, PCIe Low Profile	540-BCMWM	540-BCMXM	7TKND