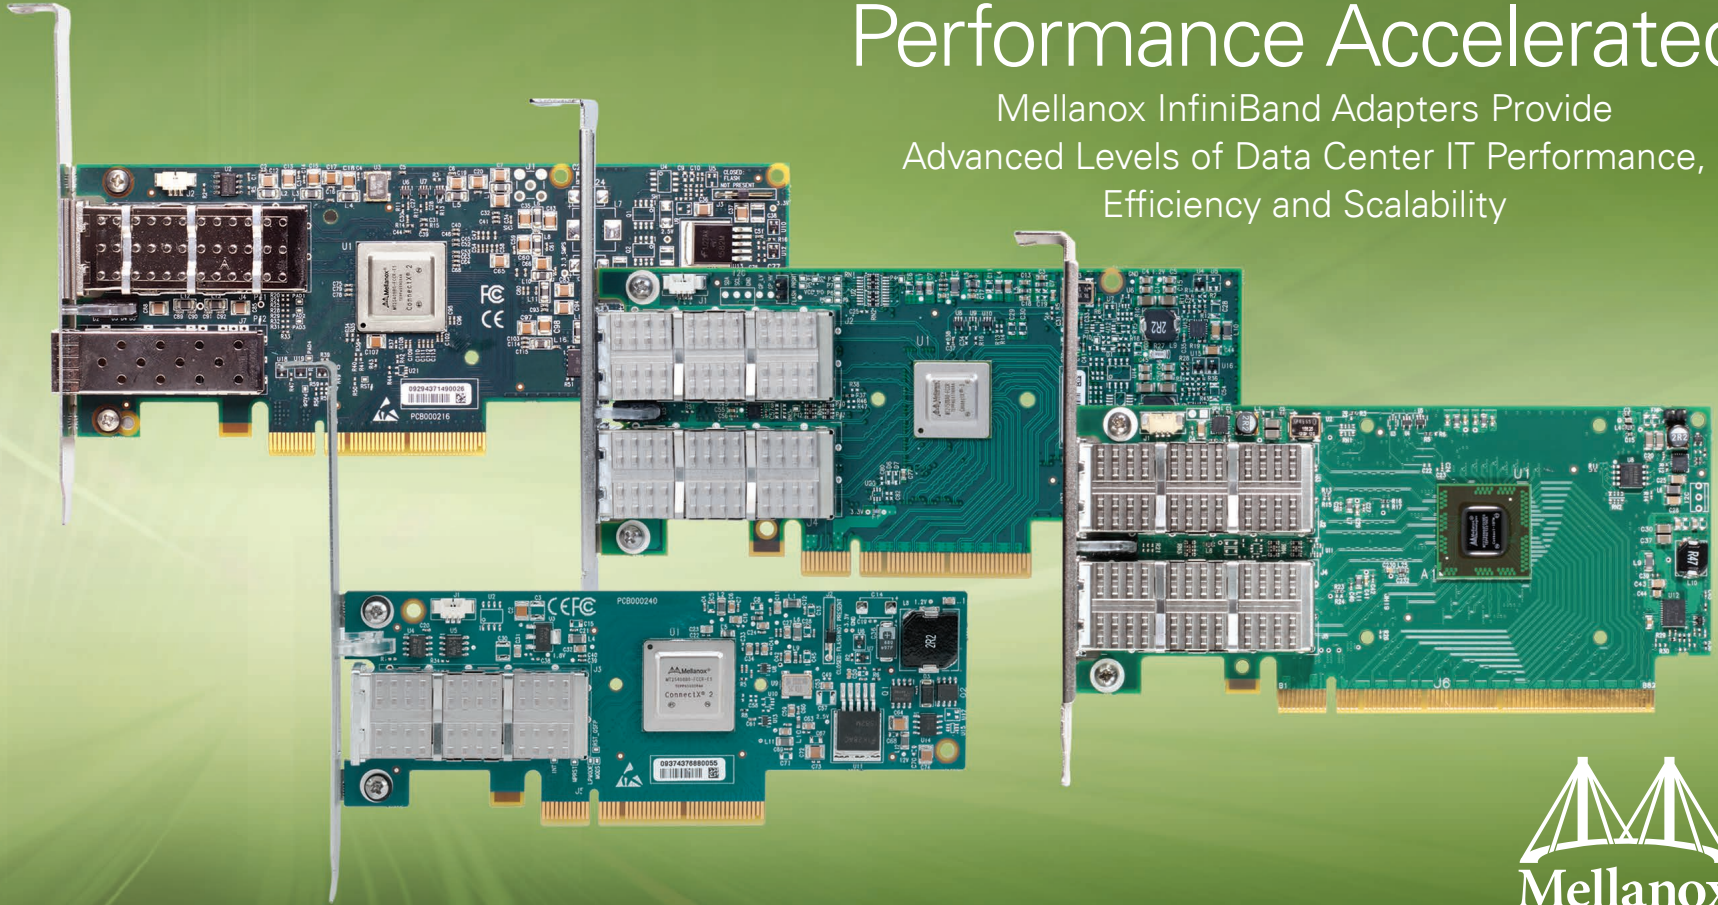


# Performance Accelerated

Mellanox InfiniBand Adapters Provide  
Advanced Levels of Data Center IT Performance,  
Efficiency and Scalability



Mellanox continues its leadership providing **InfiniBand Host Channel Adapters (HCA)** — the highest performance and most flexible interconnect solution for High-Performance Computing, Web 2.0, Cloud, data analytics, database, and storage platforms.

**Mellanox InfiniBand Host Channel Adapters (HCA)** provide the highest performing interconnect solution for High-Performance Computing, Enterprise Data Centers, Web 2.0, Cloud Computing, and embedded environments. Clustered data bases, parallelized applications, transactional services and high-performance embedded I/O applications will achieve significant performance improvements resulting in reduced completion time and lower cost per operation.

- High Performance Computing needs high bandwidth, high message rate, low latency, and CPU offloads to get the highest server efficiency and application productivity. Mellanox HCAs deliver the highest bandwidth, and message rate, and lowest latency of any standard interconnect enabling CPU efficiencies of greater than 95%.
- Data centers, high scale storage systems and cloud computing require I/O services such as high bandwidth and server utilization to achieve the maximum return on investment (ROI). Mellanox's HCAs support traffic consolidation and provides hardware acceleration for server virtualization.
- Virtual Protocol Interconnect™ (VPI) flexibility offers InfiniBand, Ethernet, and Data Center Bridging connectivity.

## World-Class Scale

With advanced performance throughput outstanding message rate capabilities, and the new memory-saving Dynamically Connected Transport (DCT) service, ConnectX-4 and Connect-IB are poised to solve the interconnect challenges of today's and tomorrow's toughest clustered computing requirements. The architecture is built from the ground up to remove bottlenecks and provide a scaleable interconnect for the largest sized and most demanding clusters.

## World-Class Performance

Mellanox InfiniBand adapters deliver industry-leading bandwidth with ultra low-latency and efficient computing for performance-driven server and storage clustering applications. Network protocol processing and data movement overhead such as RDMA and Send/Receive semantics are completed in the adapter without CPU intervention. Application acceleration and GPU communication acceleration brings further levels of performance improvement. Mellanox InfiniBand adapters' advanced acceleration technology enables higher cluster efficiency and large scalability to tens of thousands of nodes.

## BENEFITS

- *World-class cluster performance*
- *High-performance networking and storage access*
- *Efficient use of compute resources*
- *Guaranteed bandwidth and low-latency services*
- *Smart interconnect for x86, Power, ARM, and GPU-based compute and storage platforms*
- *Cutting-edge performance in virtualized overlay networks (VXLAN and NVGRE)*
- *Increased VM per server ratio*
- *I/O unification*
- *Virtualization acceleration*
- *Scalability to tens-of-thousands of nodes*

## TARGET APPLICATIONS

- *High-performance parallelized computing*
- *Data center virtualization*
- *Public and private clouds*
- *Large scale Web 2.0 and data analysis applications*
- *Clustered database applications, parallel RDBMS queries, high-throughput data warehousing*
- *Latency sensitive applications such as financial analysis and trading*
- *Cloud and grid computing data centers*
- *Performance storage applications such as backup, restore, mirroring, etc.*

## I/O Virtualization

Mellanox adapters provide comprehensive support for virtualized datacenters with Single-Root I/O Virtualization (SR-IOV) allowing dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server. I/O virtualization on InfiniBand gives data center managers better server utilization and LAN and SAN unification while reducing cost, power, and cable complexity.

## Most Efficient Clouds

Mellanox adapters are a major component in Mellanox CloudX architecture. Mellanox adapters utilizing Virtual Intelligent Queuing (Virtual-IQ) technology with SR-IOV, provides dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server. I/O virtualization on Ethernet and InfiniBand gives data center managers better server utilization and LAN and SAN unification while reducing cost, power, and cable complexity.

Overlay networks offload and encap/decap (for VXLAN, NVGRE and Geneve) enable highest bandwidth while freeing the CPU for application tasks. Mellanox adapters enable high bandwidth and more virtual machines per server ratio.

## Storage Accelerated

A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols leveraging InfiniBand RDMA result in high-performance storage access. Mellanox adapters support SRP, iSER, NFS RDMA, SMB Direct as well as SCSI and iSCSI storage protocols. ConnectX-4 and Connect-IB also bring innovative and flexible signature handover mechanism based on advanced T-10/DIF and DIX implementation. In addition, ConnectX-4 delivers advanced Erasure Coding

offloading capability, enabling distributed RAID (Redundant Array 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.

## Coherent Accelerator Processor Interface (CAPI)

ConnectX-4 enabled CAPI provides the best performance for Power and OpenPower based platforms. Such platforms benefit from better interaction between the Power CPU and the ConnectX-4 adapter, lower latency, higher efficiency of storage access, and better Return on Investment (ROI), as more applications and more Virtual Machines run on the platform

## Software Support

All Mellanox adapters are supported by a full suite of drivers for Microsoft Windows, Linux distributions, VMware, and Citrix XENServer. The adapters support OpenFabrics-based RDMA protocols and software, and the stateless offloads are fully interoperable with standard TCP/UDP/IP stacks. The adapters are compatible with configuration and management tools from OEMs and operating system vendors.

## Virtual Protocol Interconnect

VPI® flexibility enables any standard networking, clustering, storage, and management protocol to seamlessly operate over any converged network leveraging a consolidated software stack. Each port can operate on InfiniBand, Ethernet, or Data Center Bridging (DCB) fabrics, and supports IP over InfiniBand (IPoIB), Ethernet over InfiniBand (EoIB) and RDMA over Converged Ethernet (RoCE and RoCEv2). VPI simplifies I/O system design and makes it easier for IT managers to deploy infrastructure that meets the challenges of a dynamic data center.

## ConnectX-4

ConnectX-4 adapter cards with Virtual Protocol Interconnect (VPI), supporting EDR 100Gb/s InfiniBand and 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.

ConnectX-4 provides an unmatched combination of 100Gb/s bandwidth in a single port, the lowest available latency, 150 million messages per second and specific hardware offloads, addressing both today's and the next generation's compute and storage data center demands.

## Connect-IB

Connect-IB delivers leading performance with maximum bandwidth and low latency, resulting in the highest computing efficiency for performance-driven server and storage applications. Maximum bandwidth is delivered across PCI Express 3.0 x16 and two ports of FDR InfiniBand, supplying more than 100Gb/s of throughput together with consistent low latency across all CPU cores. Connect-IB also enables PCI Express 2.0 x16 systems to take full advantage of FDR, delivering at least twice the bandwidth of existing PCIe 2.0 solutions.

## ConnectX-3 and ConnectX-3 Pro

Mellanox's industry-leading ConnectX-3 family of Virtual Protocol Interconnect (VPI) adapters provides the highest performing and most flexible InfiniBand and Ethernet interconnect solution. ConnectX-3 family delivers up to 56Gb/s throughput across the PCI Express 3.0 host bus.

These adapters enable the fastest transaction latency, less than 1 usec, and can deliver more than 90M MPI messages/second making it the most scalable and suitable solution for transaction-demanding applications. The ConnectX-3 family maximizes the network efficiency making it ideal for HPC or converged data centers operating a wide range of applications.

ConnectX-3 Pro has additional dedicated hardware offloads for virtualized overlay networks (VXLAN, NVGRE) required in Cloud (IaaS) environments. These new virtualization features enable cloud service providers to efficiently expand their data centers and the service they can offer.

## Complete End-to-End 100Gb/s InfiniBand Networking

ConnectX-4 adapters are part of Mellanox's full EDR 100Gb/s InfiniBand end-to-end portfolio for data centers and high-performance computing systems, which includes switches, application acceleration packages, and cables. Mellanox's Switch-IB family of EDR InfiniBand switches and Unified Fabric Management software incorporate advanced tools that simplify networking management and installation, and provide the needed capabilities for the highest scalability and future growth. Mellanox's HPC-X collectives, messaging, and storage acceleration packages deliver additional capabilities for the ultimate server performance, and the line of FDR copper and fiber cables ensure the highest interconnect performance. With Mellanox end to end, IT managers can be assured of the highest performance, most efficient network fabric.

ConnectX<sup>®</sup> 4




Ports	1	2	1	2
Speed	FDR (56Gb/s) and 40/56GbE	FDR (56Gb/s) and 40/56GbE	EDR (100Gb/s) and 100GbE	EDR (100Gb/s) and 100GbE
Connector	QSFP			
Host Bus	PCI Express 3.0 x16			
Features	RDMA, GPUDirect, SR-IOV, Stateless Offloads, Signature Handover, Dynamically Connected Transport			
OS Support	RHEL, CentOS, SLES, OEL, Windows, ESX/vSphere, Ubuntu, Citrix, Fedora, FreeBSD			
Ordering Number	MCX455A-F	MCX456A-F	MCX455A-E	MCX456A-E

ConnectIB<sup>™</sup>



Ports	1, 2	1, 2	1, 2
Speed	SDR, DDR, QDR, FDR	SDR, DDR, QDR, FDR	SDR, DDR, QDR, FDR
PCIe Interface	PCIe 2.0 x16	PCIe 3.0 x8	PCIe 3.0 x16
Features	Hardware-based Transport and Application Offloads, RDMA, GPU Acceleration, Dynamic Connected Transport, QoS and Congestion Control		
OS Support	RHEL, CentOS, SLES, OEL, Windows, ESX/vSphere, Ubuntu, MRG, Fedora		
Ordering Number		MCB191A-FCAT MCB192A-FBAT	MCB194A-FCAT MCB193A-FBAT





Ports	1			2		
Speed	QDR IB (40Gb/s) and 10GbE	FDR10 IB (40Gb/s) and 10GbE	FDR IB (56Gb/s) and 40GbE	QDR IB (40Gb/s) and 10GbE	FDR10 IB (40Gb/s) and 10GbE	FDR IB (56Gb/s) and 40GbE
Connector	QSFP					
Host Bus	PCI Express 3.0					
Features	VPI, Hardware-based Transport and Application Offloads, RDMA, GPU Communication Acceleration, I/O Virtualization, QoS and Congestion Control; IP Stateless Offload; Precision Time Protocol					
OS Support	RHEL, SLES, Windows, ESX					
Ordering Number	MCX353A-QCCT	MCX353A-TCBT	MCX353A-FCBT	MCX354A-QCCT	MCX354A-TCBT	MCX354A-FCBT



Ports				
Speed	FDR10 IB (40Gb/s) and 10GbE	FDR IB (56Gb/s) and 40GbE	FDR10 IB (40Gb/s) and 10GbE	FDR IB (56Gb/s) and 40GbE
Connector	QSFP			
Host Bus	PCI Express 3.0			
Features	VPI, Hardware-based Transport, Virtualization and Application Offloads, RDMA, GPU Communication Acceleration, I/O Virtualization, QoS and Congestion Control; IP Stateless Offload; Precision Time Protocol			
OS Support	RHEL, SLES, Windows, ESX			
Ordering Number	MCX353A-TCCT	MCX353A-FCCT	MCX354A-TCCT	MCX354A-FCCT

## FEATURE SUMMARY\*

### INFINIBAND

- IBTA Specification 1.3 compliant
- 10, 20, 40, 56 or 100Gb/s per port
- RDMA, Send/Receive semantics
- Hardware-based congestion control
- Atomic operations 16 million I/O channels
- 9 virtual lanes: 8 data + 1 management

### ENHANCED FEATURES

- Hardware-based reliable transport
- Collective operations offloads
- GPU communication acceleration
- Hardware-based reliable multicast
- Extended Reliable Connected transport
- Enhanced Atomic operations

### ADDITIONAL CPU OFFLOADS

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

### HARDWARE-BASED I/O VIRTUALIZATION

- Single Root IOV
- Address translation and protection
- Multiple queues per virtual machine
- VMware NetQueue support

### OVERLAY NETWORKS

- Hardware offload of encapsulation and decapsulation of NVGRE and VXLAN overlay networks
- VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks
- NVGRE: Network Virtualization using Generic Routing Encapsulation

### STORAGE OFFLOADS

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

### FLEXBOOT™ TECHNOLOGY

- Remote boot over InfiniBand
- Remote boot over Ethernet
- Remote boot over iSCSI

## COMPLIANCE

### SAFETY

- UL 60950-1
- CAN/CSA-C22.2 No. 60950-1
- EN 60950-1
- IEC 60950-1

### EMC (EMISSIONS)

- FCC Part 15 (CFR 47) ,Class A
- ICES-003 ,Class A
- EN55022 ,Class A
- CISPR22 ,Class A
- AS/NZS CISPR 22, Class A ( RCM mark)
- VCCI Class A
- EN55024
- KC ( Korea)

### ENVIRONMENTAL

- EU: IEC 60068-2-64: Random Vibration
- EU: IEC 60068-2-29: Shocks, Type I / II
- EU: IEC 60068-2-32: Fall Test

### OPERATING CONDITIONS

- Operating temperature: 0 to 55° C
- Air flow: 100LFM @ 55° C
- Requires 3.3V, 12V supplies

## 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
- Coherent Accelerator Processor Interface (CAPI)

### CONNECTIVITY

- Interoperable with InfiniBand or 10/40GbE switches
- Passive copper cable
- Powered connectors for optical & active cable support
- QSFP to SFP+ connectivity through QSA module

### OPERATING SYSTEMS/DISTRIBUTIONS

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

### PROTOCOL SUPPORT

- OpenMPI, IBM PE, OSU MPI (MVAPICH/2), Intel MPI, Platform MPI, UPC, Mellanox SHMEM
- TCP/UDP, EoB, IPoB, SDP, RDS
- SRP, iSER, NFS RDMA , SMB Direct
- uDAPL

\*This brief describes hardware features and capabilities. Please refer to the driver release notes on [mellanox.com](http://mellanox.com) for feature availability.

\* Product images may not include heat sync assembly; actual product may differ.

© Copyright 2014. Mellanox Technologies. All rights reserved.

Mellanox, BridgeX, ConnectX, CORE-Direct, InfiniBridge, InfiniHost, InfiniScale, MLNX-OS, PhyX, SwitchX, UFM, Virtual Protocol Interconnect and Voltaire are registered trademarks of Mellanox Technologies, Ltd. Connect-IB, CoolBox, FabricIT, Mellanox Federal Systems, Mellanox Software Defined Storage, MetroX, Open Ethernet, ScalableHPC and Unbreakable-Link are trademarks of Mellanox Technologies, Ltd. All other trademarks are property of their respective owners.



350 Oakmead Parkway, Suite 100  
Sunnyvale, CA 94085

Tel: 408-970-3400  
Fax: 408-970-3403  
[www.mellanox.com](http://www.mellanox.com)