



ConnectX[®]-3 Pro

Single/Dual-Port Adapter Silicon with *Virtual Protocol Interconnect[®]*

ConnectX-3 Pro Adapter Silicons with Virtual Protocol Interconnect (VPI), supporting InfiniBand and Ethernet connectivity with hardware offload engines to Overlay Networks (“Tunneling”), provide the highest performing and most flexible interconnect solution for PCI Express Gen3 servers used in public and private clouds, Enterprise Data Centers, and High Performance Computing (HPC).



Public and private cloud clustered databases, parallel processing, transactional services, and high-performance embedded I/O applications will achieve significant performance improvements resulting in reduced completion time and lower cost per operation. ConnectX-3 Pro with VPI also simplifies system development by serving multiple fabrics with one hardware design.

Virtual Protocol Interconnect

VPI-enabled adapters enable any standard networking, clustering, storage, or management protocol to seamlessly operate over any converged network leveraging a consolidated software stack. With auto-sense capability, each ConnectX-3 Pro VPI port can identify and operate on InfiniBand, Ethernet, or Data Center Bridging (DCB) fabrics. FlexBoot™ provides additional flexibility by enabling servers to boot from remote InfiniBand or LAN storage targets. ConnectX-3 Pro with VPI and FlexBoot simplifies I/O system design and makes it easier for IT managers to deploy infrastructure that meets the challenges of a dynamic data center.

World-Class Performance

Virtualized Overlay Networks —

Infrastructure as a Service (IaaS) cloud demands that data centers host and serve multiple tenants, each with their own isolated network domain over a shared network infrastructure. To achieve

maximum efficiency, data center operators are creating overlay networks that carry traffic from individual Virtual Machines (VMs) in encapsulated formats such as Network Virtualization using Generic Routing Encapsulation (NVGRE) and Virtual Extensible Local Area Network (VXLAN) over a logical “tunnel,” thereby decoupling the workload’s location from its network address. Overlay Networks architecture introduces an additional layer of packet processing at the hypervisor level, adding and removing protocol headers for the encapsulated traffic. The new encapsulation prevents many of the traditional “offloading” capabilities (e.g. checksum, Time Sharing Option {TSO}) from being performed at the NIC.

ConnectX-3 Pro effectively addresses the increasing demand for an overlay network, enabling superior performance by introducing advanced NVGRE and VXLAN hardware offload engines that enable the traditional offloads to be performed on the encapsulated traffic. With ConnectX-3 Pro, data center operators can decouple the overlay network layer from the physical NIC performance, thus achieving native performance in the new network architecture.

I/O Virtualization — ConnectX-3 Pro SR-IOV technology provides dedicated adapter resources and guaranteed isolation and protection for

HIGHLIGHTS

BENEFITS

- One design for InfiniBand, Ethernet (10/40/56GbE), or Data Center Bridging fabrics
- World-class cluster, network, and storage performance
- Cutting edge performance in virtualized overlay networks (VXLAN and NVGRE)
- Guaranteed bandwidth and low-latency services
- I/O consolidation
- Virtualization acceleration
- Power efficient
- Scales to tens-of-thousands of nodes

KEY FEATURES

- Virtual Protocol Interconnect
- 1us MPI ping latency
- Up to 56Gb/s InfiniBand or 56 Gigabit Ethernet per port
- Single- and Dual-Port options available
- PCI Express 3.0 (up to 8GT/s)
- CPU offload of transport operations
- Application offload
- GPU communication acceleration
- Precision Clock Synchronization
- HW Offloads for NVGRE and VXLAN encapsulated traffic
- End-to-end QoS and congestion control
- Hardware-based I/O virtualization
- Ethernet encapsulation (EoIB)
- 17mm x 17mm RoHS-R6

virtual machines (VMs) within the server. I/O virtualization with ConnectX-3 Pro gives data center managers better server utilization while reducing cost, power, and cable complexity.

InfiniBand — ConnectX-3 Pro delivers low latency, high bandwidth, and computing efficiency for performance-driven server and storage clustering applications. Efficient computing is achieved by offloading from the CPU protocol processing and data movement overhead such as Remote Direct Memory (RDMA) Access and Send/Receive semantics, allowing more processor power for the application. CORE-Direct™ brings the next level of performance improvement by offloading application overhead such as data broadcasting and gathering, as well as global synchronization communication routines. GPU communication acceleration provides additional efficiencies by

eliminating unnecessary internal data copies to significantly reduce application run time. ConnectX-3 Pro advanced acceleration technology enables higher cluster efficiency and large scalability to tens of thousands of nodes.

RDMA over Converged Ethernet (RoCE)

- ConnectX-3 Pro utilizing IBTA RoCE technology delivers similar low-latency and high-performance over Ethernet networks. Leveraging Data Center Bridging capabilities, RoCE provides efficient low latency RDMA services over Layer 2 Ethernet. With link-level interoperability in existing Ethernet infrastructure, Network Administrators can leverage existing data center fabric management solutions.

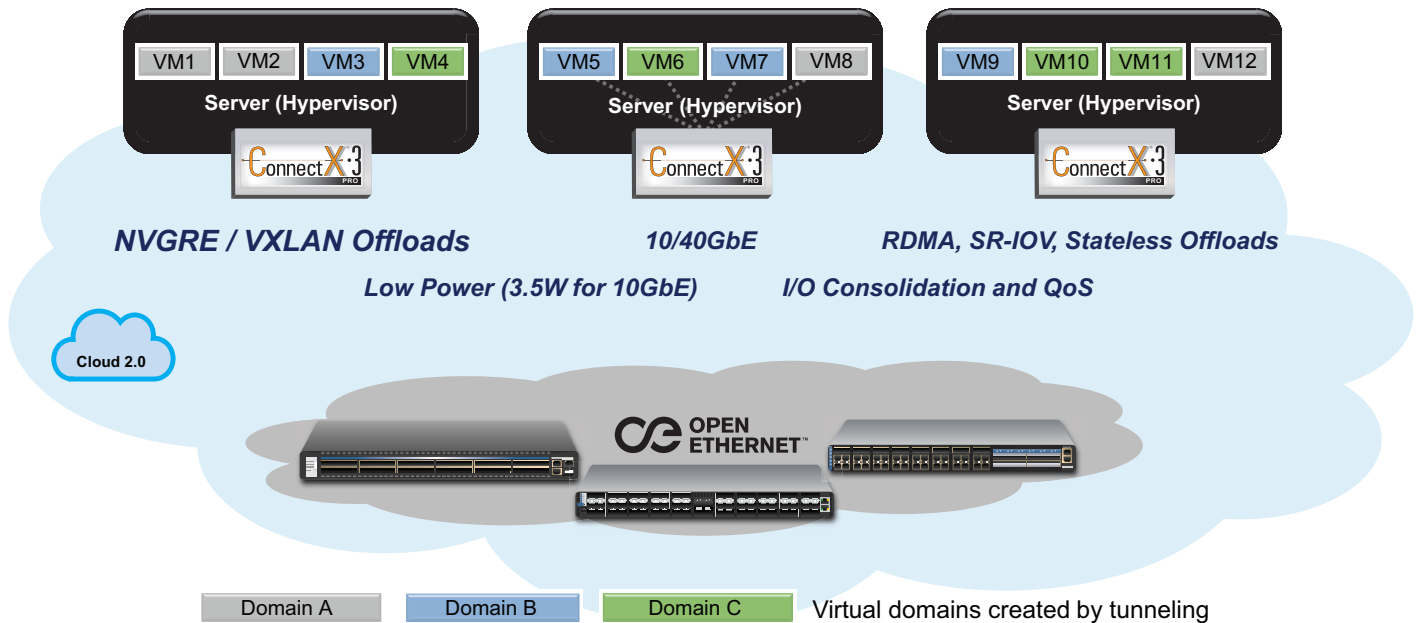
Sockets Acceleration — Applications utilizing TCP/UDP/IP transport can achieve industry-leading throughput over InfiniBand or 10/40/56GbE. The hardware-based

stateless offload engines in ConnectX-3 Pro reduce the CPU overhead of IP packet transport. Socket acceleration software further increases performance for latency sensitive applications.

Storage Acceleration — A consolidated compute and storage network achieves significant cost-performance advantages over multi-fabric networks. Standard block and file access protocols can leverage Ethernet or InfiniBand RDMA for high-performance storage access.

Software Support

All Mellanox adapter cards are supported by Windows, Linux distributions, VMware, FreeBSD, Ubuntu, and Citrix XenServer. ConnectX-3 Pro VPI adapters support OpenFabrics-based RDMA protocols and software and are compatible with configuration and management tools from OEMs and operating system vendors.



ConnectX-3 Pro Intelligent NIC — The Foundation of Cloud 2.0

FEATURES SUMMARY*

INFINIBAND

- IBTA Specification 1.2.1 compliant
- Hardware-based congestion control
- 16 million I/O channels
- 256 to 4Kbyte MTU, 1Gbyte messages

ENHANCED INFINIBAND

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

ETHERNET

- IEEE Std 802.3ae 10 Gigabit Ethernet
- IEEE Std 802.3ba 40 Gigabit Ethernet
- IEEE Std 802.3ad Link Aggregation
- IEEE Std 802.3az Energy Efficient Ethernet
- IEEE Std 802.1Q, .1P VLAN tags and priority
- IEEE Std 802.1Qau Congestion Notification
- IEEE Std 802.1Qbg
- IEEE P802.1Qaz D0.2 ETS
- IEEE P802.1Qbb D1.0 Priority-based Flow Control
- IEEE 1588v2
- Jumbo frame support (9600B)

OVERLAY NETWORKS

- VXLAN and NVGRE - A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks. Network Virtualization hardware offload engines

HARDWARE-BASED I/O VIRTUALIZATION

- Single Root IOV
- Address translation and protection
- Dedicated adapter resources
- Multiple queues per virtual machine
- Enhanced QoS for vNICs
- VMware NetQueue support

ADDITIONAL CPU OFFLOADS

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

FLEXBOOT™ TECHNOLOGY

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

PROTOCOL SUPPORT

- Open MPI, OSU MvAPIch, Intel MPI, MS MPI, Platform MPI
- TCP/UDP, EoIB, iPoIB, RDS
- SRP, iSER, NFS RDMA
- uDAPL

MANAGEMENT AND CONTROL

- NCSI, SMBus - Baseboard Management Controller interface
- I2C interface for device control and configuration
- MDIO I/F for external PHY configuration & monitoring
- General Purpose I/O pins
- SPI interface to Flash

COMPATIBILITY

PCI EXPRESS INTERFACE

- PCIe Base 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

CONNECTIVITY

- Interoperable with InfiniBand or 10/40GbE Ethernet switches. Interoperable with 56GbE Mellanox Switches.
- Passive copper cable with ESD protection
- Powered connectors for optical and active cable support
- QSFP to SFP+ connectivity through QSA module

OPERATING SYSTEMS/DISTRIBUTIONS

- Citrix XenServer 6.1
- RHEL/CentOS 5.X and 6.X, Novell SLES10 SP4; SLES11 SP1, SLES 11 SP2, OEL, Fedora 14,15,17, Ubuntu 12.04
- Windows Server 2008/2012/2012 R2
- FreeBSD
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF)
- VMware ESXi 4.x and 5.x

*This brief describes hardware features and capabilities. Please refer to the driver release notes on mellanox.com for feature availability.

**Image depicts sample product only; actual product may differ.

Ordering Part Number	Description
MT27524A1-FCCF-FV	ConnectX®-3 Pro VPI, 1-Port IC, FDR/56GbE, PCIe 3.0 8GT/s (RoHS2 R6) with HW offloads for NVGRE and VxLAN
MT27528A1-FCCF-FV	ConnectX®-3 Pro VPI, 2-Port IC, FDR/56GbE, PCIe 3.0 8GT/s (RoHS2 R6) with HW offloads for NVGRE and VxLAN



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
 Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com