Mellanox TECHNOLOGIES

Programmable ConnectX®–3 Pro Adapter Card

Dual-Port Programmable Adapter Evaluation Board

Mellanox programmable adapters provide users with the capability to program an FPGA attached to the ConnectX-3 Pro network adapter device, taking advantage of ConnectX-3 Pro enhanced application acceleration and high speed network.

The requirements of high performance and maximum flexibility in modern data centers, public and private clouds, Web 2.0 infrastructures, telecommunication, and high-performance computing are mandatory in order to achieve reduced completion time and lower cost per operation. The Programmable ConnectX-3 Pro Adapter Card simplifies system development by serving multiple fabrics with one hardware design.

Programmable Adapter - Evaluation Board

Programmable ConnectX-3 Pro adapter evaluation board supports up to 40Gb/s Ethernet connectivity with hardware offload engines. The evaluation board supports the configuration of FPGA as "bump-on-the- wire". The attached FPGA is connected between the ConnectX-3 Pro and the external switch over a 40GbE link. Moreover, additional direct PCIe Gen 3 x8 connection to the FPGA is available for high speed FPGA

configuration, code loading, or custom CPU-offload applications.

Programmable Adapters Portfolio

The Programmable Adapters portfolio includes intergrated FPGAs that can reside on the external network interfaces, on the internal PCIe bus, or on both sides. End users and partners can program the FPGA as a "bump-on-the-wire" to handle packets that are transmitted to or received from the network. The FPGA, when placed on the PCIe bus, acts as an application accelerating engine.
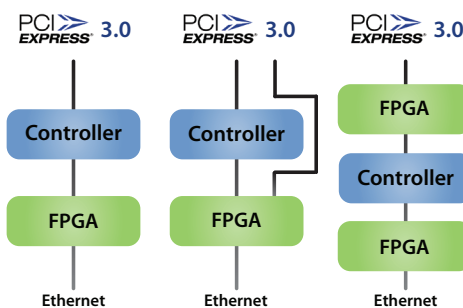
Network Adapter Logical View

HIGHLIGHTS

BENEFITS

– 10/40Gb/s FPGA as bump-on-the-wire
– FPGA on PCIe Gen3 x8 bus (up to 8GT/s) for high speed FPGA configuration
– Enabler for user application of per-packet encryption/decryption
– Enabler for CPU offload applications customization based on direct PCIe access to the FPGA
– One design for Ethernet (10/40GbE), or Data Center Bridging fabrics
– World-class cluster, network, and storage performance
– Cutting edge performance in virtualized overlay networks (VXLAN and NVGRE)
– I/O consolidation
– Virtualization acceleration
– Scales to tens-of-thousands of nodes

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 NVGRE and VXLAN over a logical "tunnel", thereby decoupling the workload location from its network address.

Overlay network 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, TSO) from being performed at the adapter.

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

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

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 RDMA over Converged Ethernet 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 adapters support OpenFabrics-based RDMA protocols and software and are compatible with configuration and management tools from OEMs and operating system vendors.

The FPGA card is supplied with an FPGA driver for Linux and Windows.

FEATURES SUMMARY*

FLEXIBLE ADAPTER FPGA

STRATIX V 5SGXA7 DEVICE FEATURES

Device Number: SGXEA7K2F40C2N

Adaptive Logic Modules (K) – 234,720
Logic Elements (K) – 622
Registers (K) – 938,880
M20K Memory Blocks – 2,560
M20K Memory (MBits) – 50

MLAB memory Logic Array Block (MBit) – 7.16

Variable-Precision DSP blocks - 256 Multipliers: (18x18) – 512

Fractional PLLs – 24
Global Clocks – 16
Regional Clocks – 92

Transceivers 14.1-Gbps - 36

PCIe (Gen3) hard IP Blocks - 4

DDR3 SDRAM x72 DIMM Interfaces – 6

Package - flipchip FineLine BGA (1.0-mm pitch): KF40-F1517

User I/Os: 696

Full Duplex LVDS: 174 14.1-Gbps

Transceivers: 36

ON-BOARD MEMORY

– Two 8GByte DDR3 – PC-1600 SODIMM
– QDR II+ 4.5MByte, 533MHz
– 512Byte of descriptors for direct access to host DDR memory by FPGA DMA

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 - 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

MAXIMUM POWER CONSUMPTION

– 113W

COMPATIBILITY

PCI EXPRESS INTERFACE

– Two PCIe Gen 3 x8 (Bifurcated PCIe x16 slot)
– PCIe Base 3.0 compliant, 1.1 and 2.0 compatible
– 2.5, 5.0, or 8.0GT/s link rate
– Auto-negotiates to x8, x4, x2, or x1 per interface
– Support for MSI/MSI-X mechanisms

CONNECTIVITY

– Interoperable with 10/40GbE Ethernet 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

– Novell SLES, Red Hat Enterprise Linux (RHEL),
– Fedora, CentOS and other Linux distributions.
– Microsoft Windows Server
– OpenFabrics Enterprise Distribution (OFED)
– OpenFabrics Windows Distribution (WinOF)
– VMware ESX Server

*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	Network Ports	Dimensions w/o Brackets
MCX366A-BCCT	Dual 10/40GbE	Full height, Full length (31.2cm x 11.115cm) 2 slots width