



Innova™ Flex 4 Lx EN Adapter Card

10/40Gb/s Ethernet adapter card with FPGA on board, single-port QSFP, and PCIe3.0 x8

Mellanox Innova adapters integrate Mellanox best-in-class network controllers with an advanced FPGA on board, introducing a true open platform for development and deployment of customized network applications while taking advantage of Mellanox network controller's hardware-based I/O isolation, unmatched scalability and efficiency.

The Mellanox Innova Flex 4 Lx EN adapter can be used to accelerate security applications, Deep Packet Inspection, compression/decompression and many other functions that require data acceleration engines to improve performance in modern data centers, private and public clouds, Web 2.0 infrastructures, telecommunications, and high-performance computing applications.

The Innova Flex 4 Lx EN adapter natively offers industry-leading technologies, such as hardware support for RDMA over Converged Ethernet, Ethernet stateless offload engines, Overlay Networks, and GPUDirect®. The addition of the FPGA on board provides users with maximum flexibility to offload CPU utilization by migrating portions of their data processing logic to the FPGA.

ARCHITECTURE

The Innova Flex 4 Lx EN adapter supports up to 40Gb/s Ethernet connectivity with hardware offload engines. The card supports the configuration of an FPGA as a "bump-in-the-wire". The attached FPGA connects between the Innova Flex 4 Lx EN adapter and the external switch over a 40GbE link. FPGA configuration and loading of custom CPU-offload applications is done via the PCIe Gen 3 x8 interface. The Innova Flex 4 Lx EN adapter card comes equipped with the Xilinx Kintex® UltraScale™ XCKU060 FPGA.

I/O VIRTUALIZATION

The Innova Flex 4 Lx EN adapter uses SR-IOV technology to provide dedicated adapter resources and guaranteed isolation and protection for virtual machines (VMs) within the server. I/O virtualization 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. The Innova Flex 4 Lx EN adapter effectively addresses this by providing advanced NVGRE, VXLAN and GENEVE hardware offloading engines that

HIGHLIGHTS

BENEFITS

- Open platform for easy development of user applications and deployment to the FPGA
- Industry leading throughput and latency for Web 2.0, Cloud and Big Data applications
- Smart interconnect for x86, Power, Arm, and GPU-based compute and storage platforms
- Cutting-edge performance in virtualized overlay networks
- Efficient I/O consolidation, lowering data center costs and complexity
- Virtualization acceleration
- Power efficiency

FEATURES

- 10/40Gb/s FPGA as 'bump-in-the-wire'
- Erasure Coding offload
- Virtualization
- Low latency RDMA over Converged Ethernet (RoCE)
- 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

encapsulate and de-capsulate the overlay protocol headers, enabling the traditional offloads to be performed on the encapsulated traffic for these and other tunneling protocols (MPLS, QinQ, and so on). With the Innova Flex 4 Lx EN adapter, data center operators can achieve native performance in the new network architecture.

RDMA OVER CONVERGED ETHERNET (ROCE)

The Innova Flex 4 Lx EN adapter supports RoCE specifications delivering low-latency and high-performance over Ethernet networks. Leveraging data center bridging (DCB) capabilities as well as 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. The Innova Flex 4 Lx EN adapter offers advanced acceleration technology that enables higher cluster efficiency and scalability to tens of thousands of nodes.

STORAGE ACCELERATION

Storage applications will see improved performance with the higher bandwidth that the Innova Flex 4 Lx EN adapter 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.

DISTRIBUTED RAID

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

DEVELOPMENT ENVIRONMENT

Applications can be easily developed and deployed to the on-board FPGA utilizing Mellanox tools and the Xilinx standard development environment. The Innova Flex 4 Lx EN adapter card includes the FPGA's Board Support Packages (BSP), providing programmers with the necessary infrastructure to implement their own applications and reduce time-to-market.

SOFTWARE SUPPORT

All Mellanox adapter cards are supported by Windows, Linux distributions, VMware, FreeBSD, and Citrix XENServer. The Innova Flex 4 Lx EN adapter supports various management interfaces and has a rich set of tools for configuration and management of both the Innova Flex 4 Lx EN network controller and Xilinx FPGA across operating systems.

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

Operating Systems/Distributions*

- RHEL/CentOS
- Windows
- FreeBSD
- VMware
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF-2)

Connectivity

- Interoperable with 10/40Gb Ethernet switches
- Passive copper cable with ESD protection
- Powered connectors for optical and active cable support

* Not all operating systems will be supported for the first release of this card.

FEATURES

Ethernet Controller

- ConnectX-4 Lx EN

Application Accelerator FPGA

- Xilinx Kintex® UltraScale™ XCKU060

On-board Memory

- 2GByte DDR4-1600

Maximum Power Consumption

- 30W with XCKU060

Ethernet

- IEEE Std 802.3ae 10 Gigabit Ethernet
- IEEE Std 802.3ba 40 Gigabit Ethernet
- IEEE Std 802.3ad Link Aggregation
- 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)

Enhanced Features

- Hardware-based reliable transport
- Collective operations offloads
- Vector collective operations offloads
- PeerDirect™ RDMA (aka GPUDirect communication acceleration)
- 64/66 encoding
- Hardware-based reliable multicast
- 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

Hardware-Based I/O Virtualization

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

Virtualization

- SR-IOV: Up to 512 Virtual Functions
- SR-IOV: Up to 16 Physical Functions per host
 - 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

Overlay Networks

- RAID offload - erasure coding (Reed-Solomon) offload
- Stateless offloads for overlay networks and tunneling protocols
- Hardware offload of encapsulation and decapsulation of NVGRE and VXLAN overlay networks

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

Remote Boot

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

* 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

OPN	Description	Dimensions w/o Bracket
MNV101512A-BCAT	Mellanox Innova Flex 4 Lx EN, single port QSFP, 40GbE, PCIe3.0 x8, active cooling	Half height, half length (68.9mm x 167.65mm)