



Innova™ IPsec Adapter Card

10/40Gb/s Ethernet adapter card delivering high-performance inline IPsec encryption acceleration combined with advanced network offloads for security at every endpoint

Mellanox Innova IPsec Adapter Card accelerates IPsec cryptographic functions at wire speed to enable a more efficient use of compute resources for the most demanding cloud, Web 2.0, telecommunication and storage systems, and other applications.

Mellanox Innova IPsec Adapter Card delivers 40GbE IPsec traffic with lower CPU utilization, freeing CPU resources for application execution.

Growing concerns over Internet traffic interception by government agencies and how unencrypted information can be gathered and used have kindled a global desire for protecting privacy. As such, encryption to protect data-at-rest and data-in-motion is gaining momentum within data centers.

Encryption of data-in-motion is particularly CPU intensive. As a result, more CPU resources are being used to perform encryption functions instead of focusing on applications. Furthermore, server CPUs cannot scale to support the ever-growing volume and velocity of traffic to be processed.

Innova IPsec adapter lowers the Total Cost of Ownership (TCO) compared to discrete encryption acceleration solutions by combining encryption offload with advanced network capabilities on a single adapter.

IPSEC OFFLOAD

The Innova IPsec adapter uses FPGA based AES-GCM and AES-CBC cryptographic engines to efficiently offload IPsec compute intensive encryption and authentication tasks from the CPU and freeing it for business application execution.

The FPGA is a 'bump-in-the-wire' architecture with the encryption and decryption being performed inline with the network flow. This means that the on-board ConnectX-4 Lx adapter's advanced offloads (overlay networks, RoCE) are being maintained while the IPsec encryption feature is activated, enabling the offload of network features with encrypted packets.

I/O VIRTUALIZATION

Innova IPsec supports SR-IOV technology and enables dedicated adapter resources and guaranteed isolation and protection for virtual machines (VMs) within the server. This I/O virtualization provides data center administrators with better server utilization while reducing cost, power, and cable complexity, allowing more virtual machines and more tenants on the same hardware.

HIGHLIGHTS

BENEFITS

- Up to 4X lower CPU utilization compared to non-offloaded encryption solutions
- Minimal changes to a customer's software stack
- Maintains network adapter offloads with encrypted traffic
- Industry-leading throughput for Web 2.0, storage and telecommunications applications
- Reduced TCO due to a combination of an encryption offload engine and a network adapter on one board
- Cutting-edge performance in virtualized overlay networks

FEATURES

- AES-GCM, AES-CBC encryption/decryption and authentication algorithm offloads
- Erasure Coding offload
- Virtualization
- Low latency RDMA over Converged Ethernet
- 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

ACCELERATION FOR 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 or 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 IPsec adapter effectively addresses this by providing advanced NVGRE, VXLAN, and GENEVE hardware offloading engines that encapsulate and de-encapsulate the overlay protocol headers, enabling the traditional offloads to be performed on the encapsulated traffic for these and other tunneling protocols (GENEVE, MPLS, QinQ, and so on). With the Innova IPsec adapter, data center operators can achieve native performance in the new network architecture.

RDMA OVER CONVERGED ETHERNET (ROCE)

The Innova IPsec adapter supports RoCE specifications delivering low-latency and high-performance over Ethernet networks. Leveraging data center bridging (DCB) capabilities as well as the Innova IPsec adapter's 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 IPsec adapter's advanced acceleration technology 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 IPsec 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.

Innova IPsec also offers 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. Innova IPsec's Reed-Solomon capability introduces redundant block calculations, which, together with RDMA, achieves high performance and reliable storage access.

SOFTWARE SUPPORT

Innova IPsec adapter is supported by the Mellanox standard OFED release that includes kernel and DPDK implementations.

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 10/40GbE 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

Maximum Power Consumption

- Approx. 25W (for 40GbE traffic)

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

- 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

Security Offloads**

- IPsec offload for Linux
- IPsec offload for Windows
- Authentication algorithms: SHA-1, SHA-2
- Encryption algorithms: AES-GCM, AES-CBC (key lengths 128/256)

Storage Offloads

- RAID offload – erasure coding (Reed-Solomon) offload

Hardware-Based I/O Virtualization

- Single Root IOV
- Multi-function per port
- 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

- 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.

** Additional algorithms can be added based on business needs.

Table 1 - Part Numbers and Descriptions

OPN	Description	Dimensions w/o Bracket
MNV101512A-BCIT	Innova IPsec EN Adapter, single-port QSFP, 10/40GbE, with Crypto, PCIe3.0 x8, HHHL, active heat sink, tall bracket, ROHS R6	Half Height, Half Length (68.9mm x 167.65mm)
MNV101511A-BCIT	Innova IPsec EN Adapter, single-port QSFP, 10/40GbE, with Crypto, PCIe3.0 x8, HHHL, passive heat sink, tall bracket, ROHS R6	Half Height, Half Length (68.9mm x 167.65mm)



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