

7 Reasons to deploy a Mellanox ConnectX 10GbE

| | |
|--|---|
| 1.0 Introduction | 1 |
| 1.1 PCIe Gen2 | 2 |
| 1.2 PHY Integration | 2 |
| 1.3 Data Center Ethernet | 3 |
| 1.4 Fibre Channel over Ethernet (FCoE)..... | 3 |
| 1.5 Single Root- IO Virtualization (SR-IOV)..... | 4 |
| 1.6 Performance | 4 |
| 1.7 Software..... | 5 |
| 1.8 Summary | 6 |

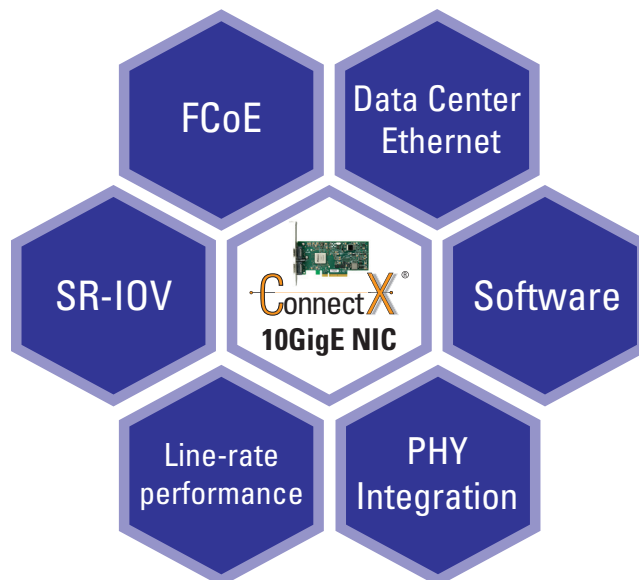
1.0 Introduction

Mellanox, the leader in the high performance interconnect technology has introduced the most compelling 10 Gigabit Ethernet silicon and adapter providing leading storage, virtualization and networking features. Mellanox dominates the 10 Gigabit Ethernet market by supporting several leading-edge technologies providing the utmost investment protection to data centers worldwide.

10 Gigabit Ethernet initially started out as a faster version of the Gigabit Ethernet technology but in the last couple of years it has become the key interconnect in the new data center. Technologies like TCP Offload, iSCSI offload and IWARP have lost significance in the new data center and new technologies like Data Center Ethernet, Fibre Channel over Ethernet, Single Root IO Virtualization have gained significant interest as these technologies future proof the networks.


Mellanox ConnectX EN 10 Gigabit Ethernet provides features that are relevant for today's deployment while ensuring investment protection and future proofing.

Mellanox ConnectX EN offers significant value by providing the following features



1.1 PCIe Gen2:

PCI Express is an industry standard for connecting IO devices. PCIe Gen2 devices are capable of supporting 5Gb/s per lane or 5GT/s compared to PCIe Gen1 devices which supports 2.5Gb/s per lane or 2.5GT/s. ConnectX EN supports a PCIe x8 link which translates to 32 Gb/s (5.0Gb/s x 8 lanes – 8b/10b encoding) per direction compared to 16Gb/s supported by PCIe Gen1 device.




ConnectX EN with Gen2 support utilizes the full bandwidth available in a Gen2 PCIe x8 slot. The dual ports in ConnectX EN not only provides high availability and failover but also link aggregation and high performance

The new servers that are being rolled out supports PCIe Gen2.

ConnectX EN delivers PCIe Gen2 on CX4 adapters, SFP+ adapters and Mezzanine card for IBM BladeCenter H

1.2 PHY Integration



The media connector for 10 Gigabit Ethernet has been evolving over the last few years. There are different types of media connectors for different types of applications. The three broad categories are: Copper, Optical and Backplane (for Blade Servers and ATCA applications).

The popular copper connector so far has been CX4 and the ConnectX EN has an integrated PHY that drives the CX4 connector. There is an emerging trend to use Twinax copper cables using SFP+ (more about SFP+ in the next paragraph)

Optical connectors were supported by using XFP technology with pluggable optics supporting SR (Short Reach) and LR (Long Reach). XFP has evolved into SFP+ which consumes less power and costs and is currently the popular media connector type as it provides the flexibility of using the optical modules for Short-reach and Long-reach applications and directly connecting copper Twinax cables for shorter distances up to 8m. SFP+ technology supports Twinax copper for top of the rack connectivity and fibre optics cables for structured cabling.

10GBASE-KR serial technology is becoming the popular PHY technology on the backplanes both on the blade server chassis as well as on the ATCA chassis. ConnectX EN integrates KR and provides a highly integrated, cost optimized and power optimized solution to both Blade server and ATCA applications.

ConnectX EN integrates all the necessary PHY technologies to deliver a low cost, low power solution.

1.3 Data Center Ethernet:

Data Center Ethernet (DCE) aka Converged Enhanced Ethernet (CEE) is part of the IEEE 802.1 Data Center Bridging Workgroup. The workgroup is focused on enhancing the Ethernet protocol to make it robust for IO consolidation in the data center. The three key enhancements that are relevant for an endpoint device like ConnectX EN are:

802.1Qbb – Priority based Flow Control aka Class based Flow Control aka Per Priority Pause. To effectively consolidate multiple traffic types on a single physical network, it is extremely important to make sure that there are flow control mechanisms to ensure problems with one traffic type does not impact the other traffic types. For example, to consolidate storage traffic over an Ethernet transport it is important that the storage traffic type does not get impacted by the network traffic types misbehavior as storage needs a reliable transport mechanism. Traffic type isolation and priority for each traffic type makes Ethernet a lossless technology. 802.1Qbb provides such an enhancement to the already existing pause technology.

802.1Qaz – Enhanced Transmission Selection: Each 802.1Qbb can manage up to eight different traffic types over a single physical wire and each of these traffic types can be further managed by providing traffic classes. Each traffic class can have its own characteristics like a low latency traffic class or the high bandwidth traffic class 802.1Qz enforces bandwidth allocation and provides the flexibility to dynamically manage the traffic and bandwidth. When traffic at a priority level does not use its full allocation, 802.1Qaz ensures that other priorities are allowed to use that bandwidth dynamically.

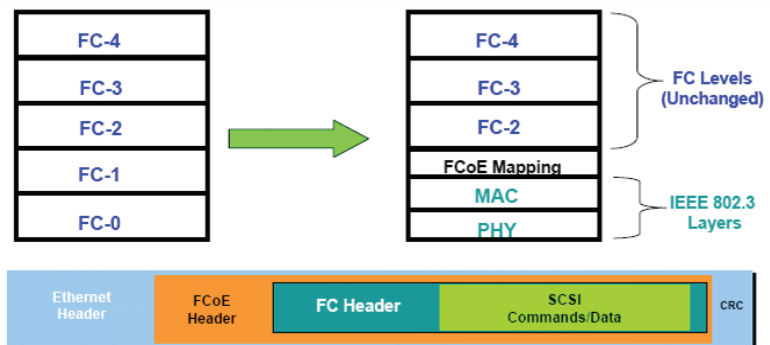
802.1Qau – Congestion Management: 802.1Qau ensures that the traffic is managed effectively at the edge of the network and reduces the effects of the network congestion. 802.1Qaz helps in identifying the congestion points and apply rate-limiting and back-pressure techniques to traffic types causing the congestion ensuring smooth traffic flow.

Mellanox ConnectX EN has built in capabilities to support all the above key enhancements and includes the support for DCBX which is a Data Center Bridging Exchange protocol for identifying DCB capable devices in the network.

1.4 Fibre Channel over Ethernet (FCoE)

During the last decade Fibre Channel technology gained significant foothold in the data center storage infrastructure. Proliferation of SANs saw the growth in FC adapters on the servers and the growth of edge FC switches. Two Ethernet ports and two FC ports per server became standard on servers deployed in the data center increasing the number of cables, power, cost and manageability. IO consolidation was imminent. The industry standards body T11 part of INCITS defined the FCoE specifications.

FCoE retains all the FC elements including the management. FCoE was defined primarily to consolidate the IO fabric on the servers to Ethernet; not the classic Ethernet that is being deployed today but Ethernet with the new Data Center Ethernet features. The most relevant to FCoE is Per Priority Pause or 802.1Qbb.



FCoE is designed to behave like a FC protocol while using Ethernet as the physical layer connectivity and allowing Ethernet switches to move the traffic ensuring IO consolidation not only on the servers but also the edge switches.

Mellanox ConnectX EN is the only product to support an offloaded FCoE solution similar to the FC

1.5 Single Root – IO Virtualization (SR-IOV):

HBAs in a single silicon and without any additional components. This results in costs and power savings while providing FC HBA like performance.

The offloaded design in ConnectX provides higher IOPS, Zero copy support and Low CPU utilization. Mellanox FCoE solution is the only solution supporting T10 DIF which is a key requirement for storage applications to maintain the data integrity.

By deploying ConnectX EN data centers can converge on a single adapter for both networking and storage traffic.


Virtualized server deployments are growing substantially in the data centers and the number of virtual machines per physical server has been increasing every year. This trend will continue because the number of cores per CPU has been doubling every 18 months which drives more virtual machines per server.

The industry consortium PCI-SIG (PCI Special Interest Group) released the specifications for virtualizing the IO using Single Root – IO Virtualization (SR-IOV). The SR-IOV is an industry standard to provide multiple virtual devices to be simultaneously shared among multiple virtual machines by partitioning a single PCIe device.

Mellanox an active participant in the PCI-SIG has incorporated the SR-IOV specifications to the ConnectX EN adapter providing the latest IO Virtualization capabilities to virtualized data centers, today.

ConnectX EN efficiently partitions its single 10Gigabit Ethernet physical function (PF) to multiple virtual functions (VF) with configurable bandwidths and memory mapping for each of the virtual functions.


ConnectX EN integrates the next generation SR-IOV technology for superior virtualization performance



1.6 Performance:

The promise of the 10 Gigabit Ethernet is 10x performance increase over the previous Ethernet standard, Gigabit Ethernet. Different applications have different performance characteristics and it is essential to ensure that performance is not characterized on only one aspect or by using only one performance tool.

High bandwidth is important for applications like video on demand, storage and back-up while latency is important for cluster interconnect. Virtualized servers need scalability as more number of virtual machines added to a physical server the IO should scale to line-rate and sustain the line rate performance.

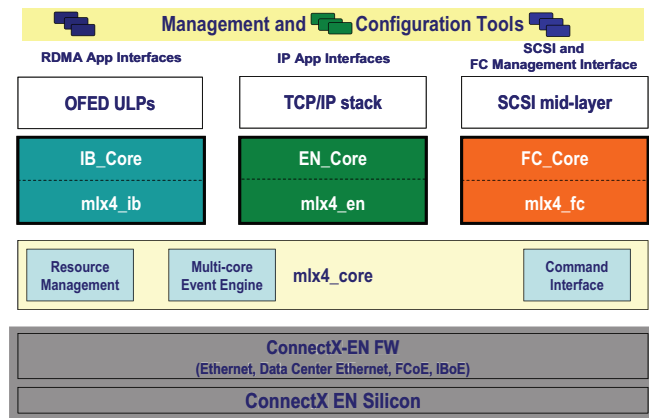


1.7 Software:

ConnectX EN in VMWare environments scales to line-rate with 3 virtual machines and sustains the line-rate performance as more number of virtual machines are added. Virtual server scaling is possible in ConnectX EN because of its ability to tie an individual TCP connection to a particular CPU core.

Similarly ConnectX EN delivers line-rate throughput with both single port and dual port in a PCIe Gen2 server providing a superior link aggregation platform.

ConnectX EN provides the best price / performance in the industry. 10 Gigabit Ethernet has evolved as the fabric of choice for IO consolidation supporting features like Data Center Ethernet, FCoE and Virtualization. Software in terms of firmware and operating system drivers play a key role in delivering a product that meets the performance expectation and robustness for an IO consolidation solution. In an IO consolidated solution it is extremely important to have an in depth understanding of all the different stacks like network, storage, virtualization, etc to provide an integrated and modular solution. Let me provide an example of a simple boot implementation. PXE boot is used for booting over a NIC card, iSCSI boot can be supported over a NIC card in the case of software initiator and FC boot using a FC HBA. In an IO consolidated environment all these three elements are possible and the management should be consistent for users to configure the appropriate boot mechanism. Now, imagine this on top of a virtualized environment where every VM could boot using a different mechanism.



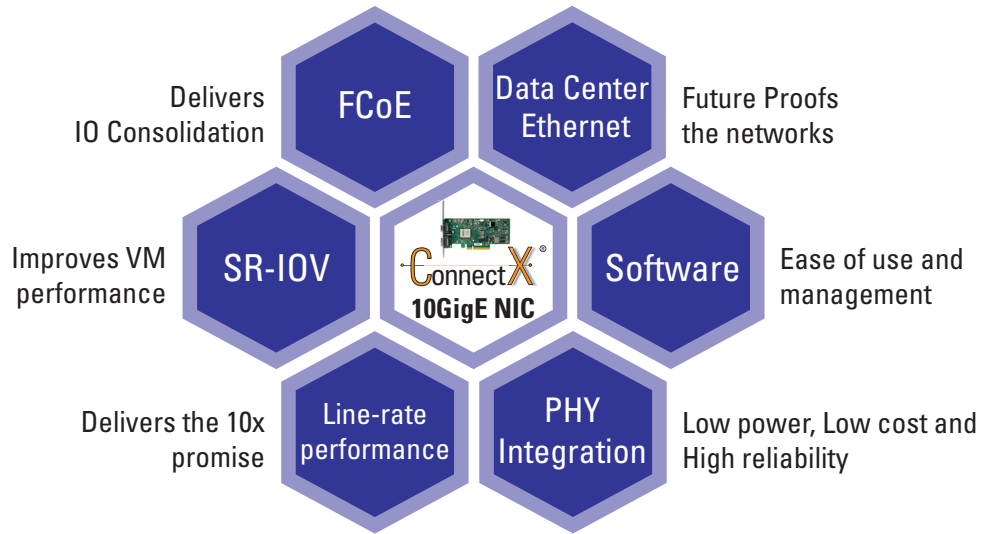
ConnectX EN provides a modular architecture and has a coherent management and configuring tools. The drivers provide rich and full functionality and are being integrated by various operating system vendors. The OS drivers maintain all application level interfaces and there are no special patches required.

1.8 Summary:

ConnectX EN delivers certified operating system drivers and rigorously tested firmware providing a fully qualified system that works the first time.

ConnectX EN is designed to deliver a complete Ethernet solution for today's and tomorrow's network by supporting Data Center Ethernet and Fibre Channel over Ethernet and other virtualization acceleration capabilities. Highly integrated hardware design provides a low CAPEX while low power utilization delivers a low OPEX. ConnectX EN leverages Mellanox experience and expertise in designing and delivering high performance interconnect for over 8 years.

ConnectX EN provides superior future proofing and investment protection while delivering high performance in both virtualized and non-virtualized operating environments



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

© Copyright 2009, Mellanox Technologies. All rights reserved. Preliminary information. Subject to change without notice. Mellanox, ConnectX, InfiniBlast, InfiniBridge, InfiniHost, InfiniRISC, InfiniScale, and InfiniPCI are registered trademarks of Mellanox Technologies, Ltd. BridgeX and Virtual Protocol Interconnect are trademarks of Mellanox Technologies, Ltd. All other trademarks are property of their respective owners.