



ARISTA: Improving Application Performance While Reducing Complexity

Problem Statement

Advances in server virtualization, network storage, and compute clusters have driven the need for faster network throughput to address application latency and availability problems in the Enterprise. Yet today's economic demands require the Enterprise maintain reasonable or reduced capital and operational cost parameters. This white paper will address how Mellanox and Arista Networks substantially improve application performance while reducing network complexity.

Problem Statement

As applications and databases increase in size, the need to have faster access to data becomes critical. Enterprises traditionally have thrown more servers to solve performance challenges. More servers bring in more I/O adapters and more I/O adapters lead to more cables and higher level of complexity as well as more power and higher cost. The number of I/O adapters in a server has been growing as each server needs connectivity to networks and to storage resulting in servers shipping with multiple, dual or quad-port Gigabit Ethernet NICs and dual-port Fibre Channel HBAs. On an average, each server in the enterprise is connected with four I/O cables increasing the complexity without addressing the real bottleneck in the data center

Enterprises are looking for improved application performance and reduced infrastructure complexity.

Previous Options: More Servers and I/O Adapters

Previously, enterprises approached performance problems by deploying additional servers and by adding I/O adapters to each server. Multiple Gigabit Ethernet NICs were installed in servers to provide more I/O leading to increased consumption of Gigabit Ethernet switch ports. This solved the problem short-term, but over time it introduced more complexity in managing the server sprawl and increased switching infrastructure and overall power consumption.

Mellanox and Arista solution

For over 8 years, Mellanox Technologies has been an active leader in high-performance interconnect technology. The Company recently introduced its feature-rich 10 Gigabit Ethernet silicon and adapter with industry leading support for networking, virtualization and storage:

Mellanox's ConnectX EN 10 Gigabit Ethernet adapter is the first adapter to support the PCI Express Gen 2.0 specification, which delivers 32Gb/s of PCI performance per direction with a x8 link compared to 16Gb/s with a Gen1 device. Dual-port Mellanox ConnectX EN adapters, not only deliver high availability, but also link aggregation and high-performance because of PCIe Gen 2.0.

Mellanox ConnectX EN 10 GbE Adapter



Different applications have different performance characteristics and Mellanox, together with their partner Arista, delivers joint solutions that address the needs of different applications with the best performance for both high-bandwidth and low-latency applications.

Mellanox provides a superior “out-of-the-box” experience by delivering a product that does not need additional tuning, configuration, nor kernel patches. Mellanox drivers leverage standard industry stacks to provide the best performance.

Mellanox and Arista conducted several performance benchmarks that provide proof points for enterprises to deploy 10 Gigabit Ethernet. Enterprises deploying Mellanox and Arista’s solutions can attain a significant performance increase over Gigabit Ethernet deployments, cabling complexity, consolidate their networking, storage and compute clusters on a single fabric and saves energy costs.

Implementation: Scenario 1: Throughput

Mellanox ConnectX EN with Arista 7124S delivers line-rate throughput with both single stream and multiple streams. Delivering line-rate on low message sizes implies that the adapter is capable of saturating the wire even in smaller packet sizes. In benchmark tests using ConnectX EN and the Arista 7124S (Figure 1, Table 1)

Figure 1. 10GbE Test Topology

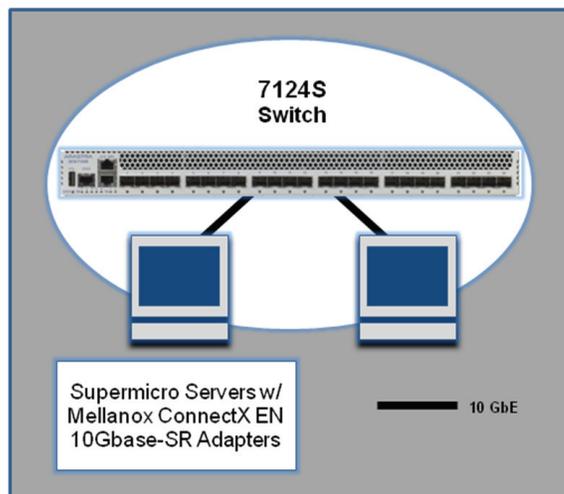


Table 1: Test Configuration
Switch: Arista 7124S w/10GBase-SR Optics
Server: Supermicro 3U X7DWE w/4GB memory
CPU: Dual Socket Quad Core Intel Xeon CPU X5482 @ 3.20GHz
Chipset: Intel Seaburg (PCIe Gen2)
OS: SLES 10 SP 1, Kernel 2.6.16.46-0.12
NIC: Mellanox ConnectX EN 10GBASE-SR w/ ConnectX: Firmware:2.5.9, Driver:MLNX_EN 1.3.0

ConnectX EN can deliver superior performance while scaling performance from message sizes of 64 Byte or 128 Byte, a size most of the enterprise applications tend to have, to large message sizes like 4 MB, which provides higher bandwidth in storage applications like back-up. Using Iperf at 1500MTU 9.47 Gbps (line-rate) was achieved. (Figure 2)

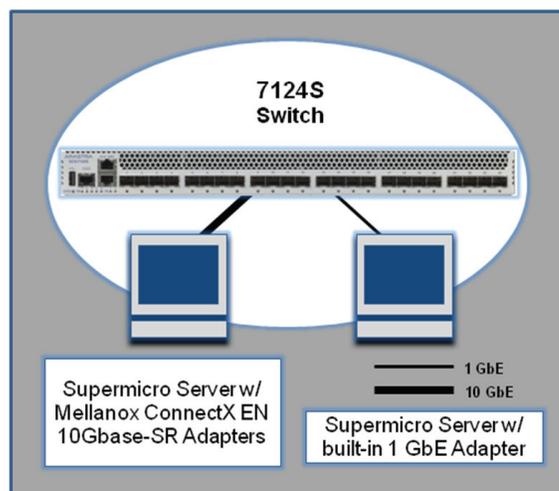
Mellanox ConnectX EN and Arista together provide the best 'out-of-the-box' performance in low latency and high bandwidth applications.

Figure 2. Test Tool: Iperf 2.0.4 Throughput Results
1500 MTU: <u>9.47Gb/s (line-rate)</u> w/message size from 128B to 4MB
9000 MTU: <u>9.92Gb/s (line-rate)</u> w/message size from 64B to 4MB

Scenario 2: TCP Latency using Standard Test Suite

The performance tests listed below uses standard operating systems without any additional tuning or running any kernel patches. These numbers are achievable in the field by using a similarly configured server with Mellanox 10 Gigabit Ethernet cards and Arista switch. (Figure 3)

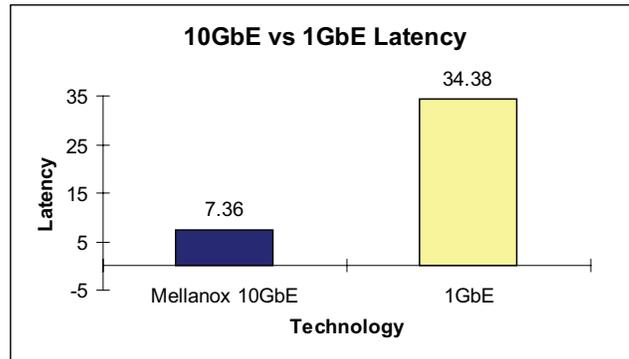
Figure 3. 10GbE vs. 1GbE Test Topology



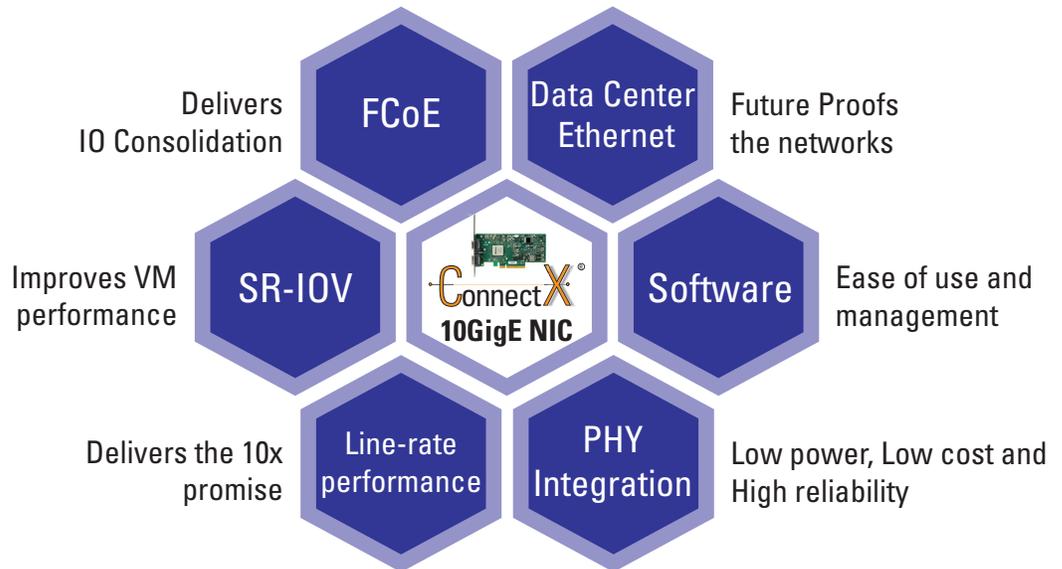
Latency on the same server using the standard built-in Gigabit Ethernet adapter was measured to be 34.38usecs. (Figure 4)

Comparing 1GbE performance versus 10GbE performance shows a 5X reduction in TCP latency by utilizing Mellanox ConnectX ED and the Arista 7124S switch.

Figure 4. 10GbE Latency vs. 1GbE Latency



Summary



Mellanox’s ConnectX EN 10 Gigabit Ethernet adapters with PCIe 2.0 delivers significant investment protection by integrating PHYs like CX4, XAU1, XFI and KR delivering low power, low cost and high reliability for today’s network and support for Data Center Ethernet, Fibre Channel over Ethernet (FCoE) and Single Root I/O Virtualization (SRIOV) for tomorrow’s networks. Enterprises can deploy the low powered and highly integrated Mellanox ConnectX EN today which provides low CAPEX while low energy costs and ease of use delivers low OPEX.

ConnectX EN leverages Mellanox’s experience and expertise in designing and delivering high-performance in both virtualized and non-virtualized operating environments. ConnectX EN supports a wide variety of operating systems like Windows Server 2003 and 2008, RedHat Enterprise Linux, Novell SuSE Enterprise Linux Server and other distributions, VMWare ESX3.5, Citrix XENSever 4.1 and FreeBSD 7.0

By deploying Mellanox ConnectX EN and Arista 7124S enterprises can get almost 5x improvements in their application performance over a 1GbE network. Low latency helps in several cluster applications like Monte Carlo simulations, portfolio analysis, compute intensive applications fast access to the data in ever increasing databases, and faster storage access. Low latency also delivers faster failover in virtualized VMware ESX environments where virtual machines can be moved from one physical server to another instantaneously]

About Mellanox:

Mellanox Technologies is a leading supplier of semiconductor-based, interconnect products to world-class server, storage, and infrastructure OEMs servicing Fortune 500 data centers, the world's most powerful supercomputers, and mission critical embedded applications. The company's 10 Gigabit Ethernet and Data Center Ethernet products with full FCoE hardware offload provide proven networking, storage, and virtualization acceleration. Founded in 1999, Mellanox Technologies is headquartered in Santa Clara, California and Yokneam, Israel. For more information, visit Mellanox at www.mellanox.com.

About Arista Networks:

Arista, based in Menlo Park, California, was founded to deliver scalable networking interconnects for large-scale datacenters and cloud computing. Arista offers best-of-breed 10 Gigabit Ethernet solutions for cloud networking that redefine scalability, robustness, and price-performance. At the core of Arista's platform is the Extensible Operating System (EOS), a pioneering new software architecture with self-healing and live in-service software upgrade capabilities. Arista's team is comprised of seasoned management and engineering talent bringing over a thousand man-years of expertise from leading network and server companies. For more information please visit: <http://www.aristanetworks.com>



2900 Stender Way, Santa Clara, CA 95054
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com

© Copyright 2008, Mellanox Technologies. All rights reserved.

Mellanox, ConnectX, InfiniBlast, InfiniBridge, InfiniHost, InfiniRISC, InfiniScale, and InfiniPCI are registered trademarks of Mellanox Technologies, Ltd. Virtual Protocol Interconnect is a trademark of Mellanox Technologies, Ltd. All other trademarks are property of their respective owners.