



Mellanox Spectrum™ and MLNX-OS based Ethernet Storage Fabric

INTRODUCTION

Storage architecture is quickly evolving to improve performance and bring cloud-like scalability and efficiency to data centers. Scale-out Storage Systems are replacing centralized systems. Fast Solid State Drives (SSDs) are replacing legacy spinning drives. As workloads move to the cloud, fewer storage systems are using legacy protocols such as Fibre Channel. Newer and more efficient protocols such as NVMe (Figure 1) are being adopted to extract even more storage performance from the infrastructure. We need fast, scalable, reliable and high-bandwidth Ethernet networks to support high-performance storage infrastructures.

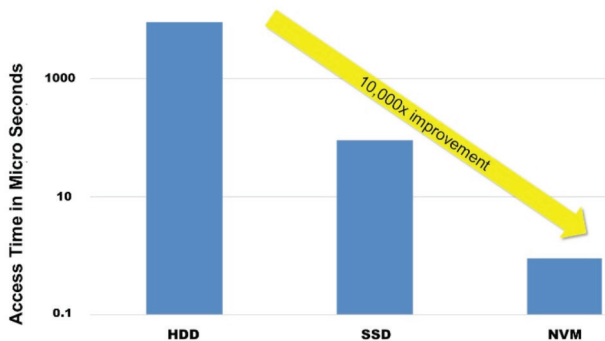


Figure 1: Dramatic Storage Performance Improvement

THE CHALLENGES WITH THE STATUS QUO

A single NVMe drive can sustain around 25Gbps of sequential read access. A server with four such drives can close to saturate a 100Gbps link. Legacy networks have performance bottlenecks and cannot sustain such high bandwidth traffic without packet drops.

NVMe protocol was specifically designed to minimize performance overheads that were part of legacy protocols such as Fibre Channel. RDMA Over Converged Ethernet (RoCE) is widely used alongside NVMe to maximize storage performance.

It is common for hyper-converged-infrastructure racks to contain 18 or fewer servers with 10GbE/25GbE connectivity. Most of the data center switches in the market today are 1U and have 48x10GbE ports. Many of the switch ports go unused due to the mismatch in the rack server density and switch port density.

SOLUTION OVERVIEW

High performance Mellanox Spectrum Ethernet switches running MLNX-OS (Figure 2) are ideal for a broad range of storage applications including NVMe-oF, Hyperconverged, Converged, scale-out or centralized storage systems. MLNX-OS makes the storage solution optimized by supporting network automation, storage visibility, and management related enhancements.





Full width (19") Switches 1 Switch fits in 1U	Half width (8.5") Switches 2 Switches fit in 1U
 <p>SN2410 Ideal 25GbE/10GbE aggregation for storage 48x10/25GbE + 8x40/100GbE</p>	 <p>SN2010 Optimized 25GbE/10GbE TOR for storage 18x10/25GbE + 4x40/100GbE</p>
 <p>SN2700 Ideal 40/100GbE aggregation for storage 32x40/100GbE or 64x50GbE or 64x10/25GbE</p>	 <p>SN2100 Ideal high-speed TOR for storage 16x40/100GbE or 32x50GbE or 64x10/25GbE</p>

Figure 2: Mellanox Spectrum Switches

KEY BENEFITS OF THE MELLANOX SPECTRUM SWITCHES

Performance

Mellanox Spectrum Ethernet switches support line rate 100GbE traffic with zero packet loss and consistent low latency. Spectrum has a robust ECN implementation which helps maximize packet throughput. The blend of high throughput, low latency and flexible form factors, makes Spectrum switches an ideal solution for storage applications.

Flexibility

Mellanox Spectrum Ethernet switches come in flexible form factors (See Figure 2). Additionally, Spectrum supports 1GbE, 10GbE, 25GbE, 40GbE, 50GbE as well as 100GbE speeds. The most optimal Spectrum platform can be chosen depending on the rack level server connectivity needs.

Spectrum Ethernet Switches can flexibly run either MLNX-OS or Cumulus Linux Network Operating System. MLNX-OS has optimizations specific to storage applications.

Efficiency

Mellanox Spectrum Ethernet switches consume low power and support inexpensive copper direct attached cables for rack-level connectivity. This results in phenomenal Opex and Capex reduction compared to legacy Fibre Channel or Ethernet solutions.

KEY BENEFITS OF MLNX-OS

Automation

MLNX-OS supports standard datacenter automation features such as ZTP, REST APIs as well as automation tools, including Ansible, Puppet, Chef, and more. In addition, MLNX-OS also supports storage-specific automation. With support for LLDP and iSCSI TLVs in MLNX-OS, storage QoS configurations can be propagated from the switch to servers.

Many high-performance storage solutions use RoCE instead of TCP/IP. MLNX-OS supports knobs to simplify end-to-end RoCE configuration.

Mellanox NEO is a powerful datacenter network orchestration platform that can be used together with MLNX-OS. NEO offers robust automation capabilities which can be used for network operations, troubleshooting and advanced visibility applications.

Visibility

MLNX-OS supports active monitoring of buffer utilization in the switch. The buffer utilization data is then summarized into compact histogram counters using hooks in the Spectrum silicon. This helps detect micro-bursts and traffic congestion issues at a very granular level. Such granular information about network congestion can be useful for debugging bursty storage applications.

Microservices

MLNX-OS supports running a container on the system. The application that runs on the container has complete access to the silicon SDK and can be used to run a customized monitoring agent that can be part of a larger monitoring and management framework. Storage specific applications such as endpoint discovery services or customized failure recovery mechanisms can now be hosted in this container infrastructure.

CONCLUSION

Mellanox Spectrum and MLNX-OS form a high performance, efficient, storage-optimized solution. Spectrum Ethernet switches support line-rate, low-latency, high-bandwidth networks with zero packet loss. MLNX-OS optimizes the solution by providing key automation, visibility, and management features that are useful for storage environments.

EXPLORE FURTHER

http://www.mellanox.com/page/ethernet_switch_overview

<http://www.mellanox.com/open-ethernet/>

<http://www.mellanox.com/tolly/>

https://www.youtube.com/Mellanox_Ethernet_Switches

About Mellanox

Mellanox Technologies (NASDAQ: MLNX) is a leading supplier of end-to-end Ethernet and InfiniBand intelligent interconnect solutions and services for servers, storage, and hyper-converged infrastructure. Mellanox intelligent interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance. Mellanox offers a choice of high performance solutions: network and multi-core processors, network adapters, switches, cables, software and silicon, that accelerate application runtime and maximize business results for a wide range of markets including high performance computing, enterprise data centers, Web 2.0, cloud, storage, network security, AI, telecom and financial services. More information is available at www.mellanox.com.



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085

Tel: 408-970-3400 • Fax: 408-970-3403

www.mellanox.com