

Faster Interconnects for Next-Generation Data Centers

Summary

With the data deluge washing over today's data centers, IT infrastructure benefits from faster interconnects. Faster storage requires faster networks. Even more performance can be achieved by using iSER, a maturing standard for extending iSCSI with RDMA (Remote Direct Memory Access). Using iSER, high-performing storage can be connected to fast Ethernet links via iSCSI, speeding data transfers from the network to servers and storage systems. These technologies can be used together to replace aging high-speed interconnects, such as Fibre Channel links and older Ethernet links.

Insight

Author: Jean S. Bozman

Introduction

A wave of change is sweeping over the data center – and a multi-terabyte data deluge is coming along with it. This rapid change is being generated by new sources of data coming from Cloud Computing, Mobility, Social Media and Big Data/Analytics. All of these workloads must now be supported, along with transactional processing, enterprise applications and data warehouses.

For many IT professionals, the new demands are causing a “rethink” of their data-transfer infrastructure. The Fibre Channel interconnects and slower Ethernet links are showing their age – and the aging links must be updated or replaced as quickly as possible. Otherwise, these slow links will be “bottlenecks” that slow down the business or organization that depends on that IT infrastructure.

This process of data center transformation requires thoughtful planning – but IT organizations cannot afford to wait too long, or performance of mission-critical applications and databases will suffer.

Flash Storage and Your Network

Many technologies will need to become more efficient in the process, including the storage systems that house the data – and send it to the compute engines for processing.

As seen in Figure 1, below, storage media is changing, with more data centers installing solid-state disks (SSDs) and even moving to Non-Volatile Memory Express (NVMe) technology to optimize SSD performance. Data-access times are falling dramatically with these new technologies, as the figure shows – and more capacity is being



housed in SSDs, rather than depending on spinning, mechanical hard-disk drives (HDDs) alone.

The rapid adoption of SSDs in server and storage systems is making it possible to have faster data-transfer rates (I/Os per second, or IOPS) – but the network interconnects themselves have to be up to the task.

Taking Stock of Your Data-Center Interconnects

Now is the time for IT organizations to take inventory of their current interconnects, and to evaluate which ones need to be updated to keep pace with the increasing requirements to transfer large datasets quickly and support high IOPS requirement applications.

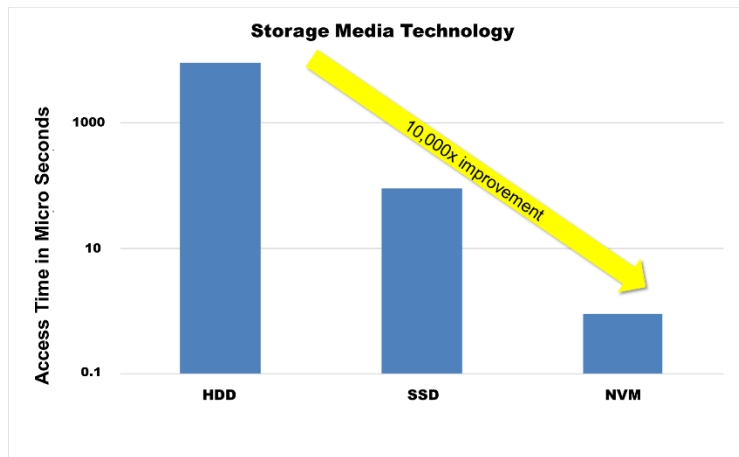


Figure 1: Improved Access Times in Storage Media Technology (Mellanox 2015)

SSDs (solid state disks) have become mainstream now, positioned between DRAM memory and HDDs (hard disk drives). It should be noted that DRAM technology is faster than SSDs, which are faster than HDDs – and that DRAM is much faster than HDDs, which are based on rotating mechanical drives.

New NVMe technology is even faster than the SAS/SATA SSD technology used in many flash-enabled storage systems. Today, the fastest flash devices support data access time in microseconds, rather than milliseconds – a dramatic reduction in access time that shows a 10,000X improvement when NVM technologies are



compared with older HDD technologies.

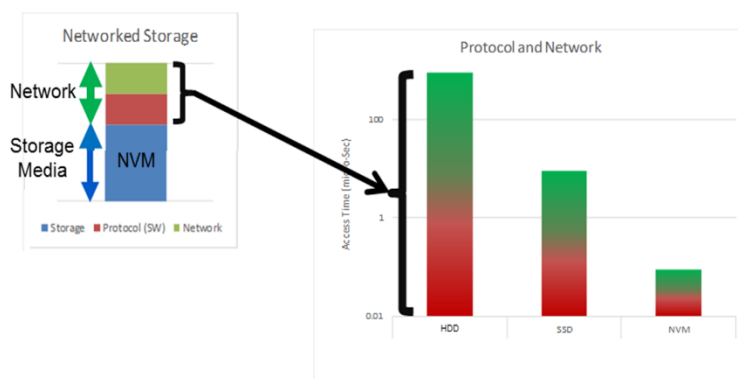


Figure 2: Improved Storage Speeds vs. Current Network Speeds (Mellanox 2015)

This improvement in storage technology moves the spotlight to the network as the next area for improvement. Network technologies to be examined include the protocol stack software, along with the wire and switching speeds of the network components.

Figure 2, above, shows the disparity between improved storage speeds, and the current options for networking speeds. This disparity prevents applications from taking full advantage of the faster storage access times now available in the data center.

To match the improvements in storage, the data center must have faster network components and faster protocol software, accelerated through the process of hardware offloads. It's worth noting here that IT organization's work to improve compute and storage speeds could end up being wasted effort when network connections speeds are not also speeded up via faster interconnects.

The Impact of Faster Ethernet iSCSI and iSER Technology

In the past, to remove the bottlenecks associated with slow networking, many organizations had been willing to pay for Fibre Channel(FC) technology for their data-center storage. While FC was very useful for providing high-speed data transfers to mission-critical systems, compute and storage is changing – with more of today's

applications relying on clusters of smaller networked servers connected by Ethernet.

Using a separate FC network for systems deployments associated with new workloads – Big Data, Cloud Computing, hyperconverged infrastructure, and software-defined storage – is typically unwieldy and, often, prohibitively expensive. Additionally, much of the installed FC infrastructure is still running at 4 or 8 Gbps speeds—or a maximum of 16 Gbps (gigabits per second), per link.

That's why many IT organizations are looking to speed up widely installed Ethernet by adding upgraded technology. Ethernet is now much faster than FC available at speeds of 10, 25, 40, 50, or 100 Gbps. New high Performance Storage can be connected to Ethernet by using iSCSI technology – an industry-standard interface that allows the SCSI data to be sent, over Ethernet local-area networks (LANs), wide-area networks (WANs) or the Internet.

It's worth noting that iSCSI protocol is natively supported on all server operating systems and most storage arrays and appliances. Upgrading the part of the network where SSD enabled systems are located will bring dramatic performance improvements to network based applications.

iSER: What Is it?

To further improve the performance of data-intensive workloads in the data center, you can use the iSER (iSCSI Extensions for RDMA) protocol on higher speed Ethernet. iSER is a new standard for extending iSCSI with Remote Direct Memory Access (RDMA).

iSER increases performance, and reduces network latency, by removing the need for the TCP/IP (Internet transport protocol), which has been in use since before the Internet existed. iSER technology was invented when SRP (SCSI RDMA Protocol) technology – originally developed for InfiniBand high-speed interconnects used in high-performance computing – was modified for Ethernet because of Ethernet's wider use throughout the world's enterprise data centers.

Figure 3 provides a view of how iSER works. Adding iSER technology increases Flash Storage performance even more than iSCSI alone by reducing the latency associated with the TCP/IP stack.



Specifically, iSER extends the iSCSI protocol to use RDMA. RDMA is provided by RoCE (RDMA over converged Ethernet), which does not need the TCP layer and therefore provides lower latency. It permits data to be transferred directly into and out of SCSI computer memory buffers (which connects computers to storage devices) without intermediate data copies.




Figure 3: How iSER Works (Mellanox 2015)

To deploy iSER in the data center, IT organizations need to support applications that use SCSI-based block storage, an operating system or hypervisor that supports an iSER initiator, an iSER-capable target – and a network that supports RDMA.

Storage Networks in Data Centers Today

Many datacenters have FC Storage Area Networks (SANs) for their mission-critical systems. These were long used to support storage for enterprise workloads in the data center. Today, as IT managers seek alternatives needed to support higher performance storage, FC interconnect are being replaced by high-speed iSCSI or iSER on Ethernet, which provide another way to support high-data-transfer rates to enterprise systems in the data center.

Figure 4, below, shows the overlap between Networking Technology choices in today's data centers. The red labels in Figure 4 show storage protocols and the purple labels in the figure show data transport options.

This sidebar outlines many of the important technical enhancements for Ethernet that have been added in recent years. Taken as a whole, these enhancements make high-performing Ethernet links a viable alternative to aging Fibre Channel interconnects.

- *VLANs and Challenge Handshake Authentication Protocol (CHAP) ensure that storage initiators and targets are isolated and that they trust each other in the same way that zoning and CHAP do on Fibre Channel networks. VLANs also enable Quality of Service (QoS) features to avoid bandwidth bottlenecks and excessive jitter in data transmission.*
- *An Ethernet protocol called Data Center Bridging provides many reliability advantages. These enable Ethernet to become a lossless network by enabling channel protocol options, which keep data off the network until there is room for it.*
- *The Server, Appliance and Array-based storage management software that runs over Fibre Channel is now supported on iSCSI and iSER. This includes the ability to boot from SAN and VMware APIs for Array Integration (VAI) primitives that enable Atomic test/set, Full copy, Block zero, Thin provisioning and other storage management features.*
- *The 40 Gbps Ethernet technology, which has been shipping since 2013, has brought performance with Ethernet to the level that was once associated with Fibre Channel. Today, the 40 Gbps Ethernet technology is running at more than twice the wire speed of 16 Gbps Fibre Channel.*

Technology for Ethernet, which is widely installed in data centers worldwide, continues to see performance improvements. With the introduction of 50 and 100Gb Ethernet upgrades in 2015, 40Gb Ethernet is no longer the fastest Ethernet on the planet. And, it's important to note that iSCSI is easily able to match Fibre Channel IOPs levels. Even greater performance can be gained with iSER, by providing more IOPs and latency reduction over iSCSI.




Figure 4: Network Technology Choices in the Data Center (Mellanox 2015)

By leveraging existing Ethernet skill-sets in working with iSCSI or iSER, data centers can achieve higher levels of standardization and interoperability between long-installed data center infrastructure and emerging new technologies.

Customers considering a move away from Fibre Channel can evaluate the enterprise-level features now offered with Ethernet [See sidebar on the left-hand side of this page]. Enhancements to the Ethernet protocol that have occurred since the early days of Fibre Channel support the high levels of performance, reliability, security, quality of service (QoS), and storage management expected for support of mission-critical enterprise workloads.

What This Looks Like in Your Data Center

Taking this approach to data-center interconnects will produce many business benefits for IT organizations, as follows:

- This approach “fits” with the technologies they already have in place – Ethernet and iSCSI. No further changes for interconnect infrastructure have to be made, avoiding a rip-and-replace project that would be disruptive to the business.
- Management software and existing IT skill-sets can be leveraged to use the faster interconnect technology, reducing Opex costs.
- Security for mission-critical data will be preserved, without the need for new and/or special data security software.
- Availability/reliability infrastructure will remain in-place, as these will all be able to leverage existing Ethernet-based infrastructure and technologies.



Mellanox Products for High Speed iSCSI and iSER

Mellanox provides a full portfolio of interconnect solutions to support highly efficient, high-performance storage networking for the iSCSI and iSER Standards. Mellanox is also expanding the ecosystem of storage partners that support its products for the iSER and iSCSI standards.

The Mellanox product portfolio includes the following:

- Adapters supporting iSER and iSCSI and are available in a variety of speeds and port configurations, ranging from 10 up to 100 Gbps Ethernet.
- Adapter software drivers are available for the most popular operating systems and hypervisors, including Linux, Microsoft Windows, VMware, and FreeBSD, with open-source iSER support in Linux, VMware, and FreeBSD.
- Switches supporting iSCSI and iSER over Ethernet, supporting lossless networking with the highest throughput, lowest latency, and lowest power consumption in the industry.
- Copper cables, optical cables and transceivers with the highest quality and lower-than-standard error rates, ensuring industry-leading support for a range of Ethernet speeds: at 10, 25, 40, 50, and 100 Gbps.

Conclusion

The data deluge has led to a rapid rise in storage capacity in the data center. To avoid storage bottlenecks for high performance applications, IT organizations are adopting solid-state flash storage – and that creates a bottleneck in the network. Faster storage needs faster networks.

These developments in the data center are causing many organizations to “take stock” of their network interconnects, identifying which links need to be updated or replaced with faster technology. Otherwise, the slow network infrastructure would prevent any gains in data-access speeds from faster storage. That’s why faster Ethernet with iSCSI or iSER is being introduced to enable solid-state flash storage performance.



IT professionals in data centers are finding that iSER leverages the compatibility, security and ease of use of familiar iSCSI-based paths to storage devices. Importantly, iSER lives in the same infrastructure that supports other types of storage and server traffic, making it widely applicable, and non-disruptive to existing operations.

By providing technologies that fit into a broad variety of networking products based on Ethernet, Mellanox is building on existing data-center infrastructure – and offering products that improve performance while avoiding rip-and-replace approaches to technology refresh.

CONTACT US

To learn more about Neuralytix and our other solutions, [contact](#) your local representative – or visit Neuralytix.com.

© Copyright 2015 Neuralytix, Inc. All Rights Reserved

The information in this publication is provided "as is." Neuralytix, Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Neuralytix believes the information contained herein is accurate as of its publication date. The information is subject to change without notice.

Neuralytix, the Neuralytix logo, the Hex logo, Neuralytix IQ are registered trademarks or trademarks of Neuralytix, Inc. All other trademarks used herein are the property of their respective owners.

