

Major Financial Institution Keeps its Data Center Efficiency One Generation Ahead



To manage the continued rapid-growth of its storage infrastructure a major financial institution decided to explore alternative to its current Fibre Channel (FC) based SAN. The company turned to Microsoft and Mellanox for an advanced storage solution using a more efficient and scalable interconnect technology that enables processing larger data sets faster and at a lower cost.

Executive Summary

The continued exponential growth in data has forced IT managers to stop deploying traditional Storage Area Network (SAN) technologies and to adopt modern architectures such as scale-out, software-defined storage, thereby enabling a higher degree of scalability and the ability to process larger amount of data faster and at a lower cost.

A major financial institution has been facing a challenge in managing the rapid-growth of its storage infrastructure, so decided to compare the option of expanding their current Fibre Channel (FC) SAN solution vs. a converged scale-out storage solution based on Microsoft's Storage Spaces using Mellanox's RDMA (Remote Direct Memory Access) enabled fabrics.

The benchmark results showed that when running Storage Spaces over SMB Direct with Mellanox 40GbE with RoCE end-to-end solution (RDMA over Converged Ethernet) the Company was able to:

- Accelerate its access to storage by 2.5 to 9 times

- Enable 8.8 times faster Virtual Machine live migration and 12.5 times faster storage migration
- Cut the number of IO ports by 4X
- Lower the storage costs by 50%

The cost savings were due to the ability of Microsoft's Storage Spaces to utilize industry-standard hardware which eliminates the cost of proprietary hardware and software required by the traditional FC-based SAN solutions.

The Situation

For more than a decade, storage architectures have remained relatively static -using SAN and NAS as the primary architectures for block and file storage, respectively. However, with recent technology trends such as Virtualization and SSD-based storage, and with the need to handle and process massive amount of data of various types in (near) real-time for mobile users, IT managers have had to evaluate and adopt new storage systems.

For example, in the "pre-virtualization" era, when each server ran one job, the access to

*"To make storage cheaper we use lots more network!
How do we make Azure Storage scale? RoCE (RDMA over Ethernet) enabled at 40GbE for Windows Azure Storage, achieving massive COGS savings"*
Albert Greenberg, Keynote at Open Networking Summit, 2014

storage was handled by only one job manipulating one type of data. It was therefore fairly easy to configure the storage system to operate efficiently. However, in a virtualized environment, where many VMs are running in parallel and each VM may require data from different locations on the disk, and when the data comes in different types (text, video, voice) and sizes (from 1B to 256KB), the traditional architectures are much less effective, as there is no way to optimize the storage systems to all of the workloads that run simultaneously.

Recently a major financial institution (Bank) decided to add a new set of internal and external services that required a massive expansion of its storage capacity and the ability to provide those services in real-time while still keeping the Total Cost of Ownership (TCO) under control. The company realized that using their current FC-based SAN architecture would be very complicated to manage and extremely expensive, and would not provide the response time required to serve their online and mobile customers. Instead, the Company decided to evaluate an advanced solution that could meet their goals at a fraction of the cost of a traditional FC SAN.

The Solution

The Bank evaluated Microsoft’s Storage Spaces solution over SMB Direct, which is Microsoft’s most recent storage solution. Storage Spaces is an integral part of Windows Server 2012 R2 and includes all the features that traditional storage provides, including Deduplication, Disaster Recovery, and advanced Quality of Service. (For additional details, please read Microsoft’s white paper “Windows Server 2012 R2 Storage”¹). This solution gives customers the ability to use low-cost industry standard hardware that they can easily scale out using SMB Direct (SMB 3.0 over RDMA²). When this solution is used in Microsoft Azure cloud, it provides double the performance (vs. the traditional SAN architecture) and 50% lower storage acquisition cost.




Figure 1. Comparison of cost acquisition between model scenarios¹

In order to compare its current FC-based SAN deployment to a Storage Spaces-based deployment, the company decided to add a file server to the existing infrastructure connected to SMB Direct using Mellanox’s end-to-end 40GbE solution, and to run several benchmarks over both configurations.

The current setup, illustrated in Figure 2, uses two different networking types. The first is 10GbE NICs and switches, and the other is 4 x 8Gb/s FC HBAs. This, of course, makes the solution more expensive, more difficult to manage, and less flexible to scale.




Figure 2. Current FC-based SAN deployment

That setup was compared to Storage Spaces running over SMB Direct, as illustrated in Figure 3, using the Mellanox end-to-end 40GbE solution that includes the 40GbE SX1012 switch, ConnectX-3 40GbE NICs, and LinkX 40GbE copper cables.




Figure 2. Deployment topology using Storage Spaces over SMB Direct

Test Results

The following graphs show the results of the various benchmarks that were run to compare the performance over the two deployments:

Benchmark	FC Based SAN	Storage Specs
Filecopy	1GB/sec	2.5GB/sec
Ntttcp	4Gb/sec	37Gb/sec
Live VM Migration	90 sec	17 sec
Live Storage Migration	327 sec	50 sec
VM Creation	13:30 min	10:45 min

These results demonstrate the superior performance that Storage Spaces provides over FC-based SAN. Storage Spaces also cut the interconnect ports from 4 to 1 and enabled the replacement of proprietary storage systems with commodity components, cutting the cost by 50%.

After reviewing these results, the major financial institution decided to proceed to the next phase in their evaluation of replacing the FC-based SAN with Microsoft's Storage Spaces over SMB Direct, using Mellanox end-to-end 40GbE RoCE.

Summary

Modern data centers need to handle massive amounts of data, and therefore must adopt new storage technologies that eliminate interconnect bottlenecks and enable higher performance and flexible scalability, all at a lower cost. Microsoft's Storage Spaces over SMB Direct delivers a scale-out storage that meets all of those needs.

Mellanox end-to-end RDMA network solutions deliver the following features:

- The fastest interconnect adapters with 10, 25, 40, 56, and 100Gb/s bandwidth per port and sub-1us latency
- I/O consolidation of server-to-server and server-to-storage over a single wire
- Full OS bypass and Remote Direct Memory Access (RDMA)
- Superior high availability solutions based on lossless interconnect technology
- Cost-effective, high-density switches and fabric architecture

Mellanox RDMA technology has been qualified by Microsoft and is fully integrated within Windows Server 2012 R2. It has been used in massive deployments such as Microsoft Azure, which results in high reliability and risk-free deployments.

Using Storage Space over RDMA solves storage bottlenecks and enables higher performance at 50% of the cost, which helps enterprises keep their data center efficiency one generation ahead.

References

¹ http://www.transtec.de/fileadmin/Medien/products/storage/Windows_Server_2012_R2_Storage_White_Paper.pdf

² http://www.mellanox.com/related-docs/applications/SB_Accelerating_Remote_Storage_Access.pdf



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