



HPE, Samsung and Mellanox Demonstrate Outstanding Efficiency of RDMA over Converged Ethernet (RoCE) with Windows Server® Storage Spaces Direct Solution

EXECUTIVE SUMMARY

Storage Spaces Direct (S2D) is a Microsoft Windows Server technology that enables you to virtualize storage by using disks and grouping them into storage pools. The capacity of storage pools can then be used to create virtual disks referred to as storage spaces which were developed to handle the exponential growth of data that is causing significant challenges within IT. Compared to traditional storage solutions, S2D can double the performance at half the cost, enabling significantly higher efficiency in Windows-based data centers.

In this paper, we'll review an industry leading S2D solution powered by technologies from four worldwide technology leaders - AMD, HPE, Samsung and Mellanox. The results demonstrate high performing and lower cost storage maximizing the storage efficiency for modern workloads. The resulting solution with RoCE demonstrates industry leading performance over conventional TCP/IP implementations.

DEMONSTRATING THE RESULTS

The performance leadership was demonstrated using a 3-nodes cluster that included HPE Cloudline CL3150 Gen10 Server, powered by a single (1P) AMD EPYC™ 7601 processor (32 cores, 64 threads and 8 memory channels in one socket).

Each server was equipped with four Samsung PM963 1.92TB NVMe SSDs, eight Samsung 32GB DDR4 DRAM with a total of 256GB memory, and one Mellanox ConnectX-5 100GbE NIC. The server was connected by HPE's StoreFabric M-Series Ethernet switch. The configuration of the test set up is depicted in figure 1.

In the tested configuration, one virtual machine (VM) was allocated per core, bringing the total numbers of Virtual Machine (VMs) in the 3-node cluster to 96. The demonstration

SOLUTION HIGHLIGHTS

- Flash Storage enables Higher Software Defined Network Efficiency
- NVMe Storage Runs Best over Higher Performance Networking Solutions
- RoCE enables up-to 5 times Higher Performance vs. TCP/IP

Figure 1 - A physical layout of our Storage Spaces Direct cluster




Figure 2 - Hardware components




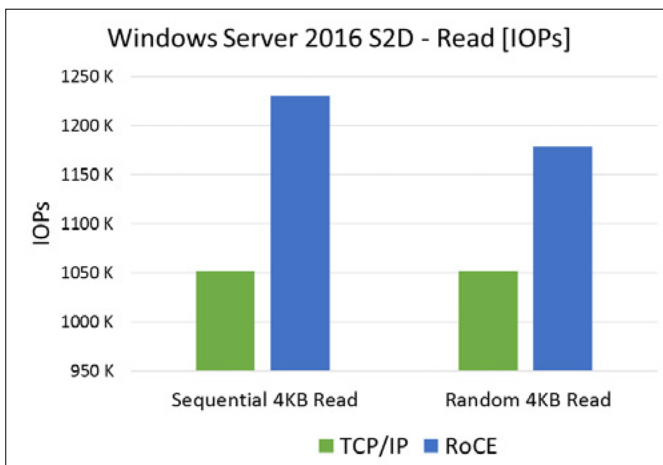
used DiskSpd to simulate application and user workloads to generate 4KB read requests from each of the 96 VMs that were controlled by Microsoft's VM Fleet benchmark (DiskSpd and VM Fleet are available as Open Source at <https://github.com/microsoft/diskspd>. DiskSpd is also available as a binary download from Microsoft at <http://aka.ms/diskspd>).

Measuring the S2D performance over the 96 VM cluster showed that read operations over RoCE generated 20% more IOPs vs. over TCP/IP. This was mainly due to the data being in cache most of the time in the TCP/IP testing, which is bypassed when using RoCE. However, in case of write operations, the number of IOPs increased to close to 5 times higher.

TESTED CONFIGURATION

The benchmark was performed over a three-node cluster and included the following components:

- HPE Cloudline CL3150 Gen 10 Server
- AMD EPYC 7601 processor (32 cores, 2.0 GHZ, 180W).
- 4 x Samsung 1.92TB NVMe SSDs (7.68GB storage)
- 8 x Samsung 32GB DDR4-2666- (256GB RAM)
- Mellanox MCX515A-CCAT 100 Gigabit Ethernet
- HPE's StoreFabric M-Series Ethernet switch



ABOUT HPE CLOUDLINE CL3150

The [HPE Cloudline CL3150 Gen10 Server](#) is a 1U 1P AMD EPYC, powered by AMD EPYC 7000 series processors, open standards-based, ultra-dense storage server. It is ideal for fast storage and real-time analytics, 3D modelling, FSI trading, rich media streaming, and e-commerce affinity advertising workloads. It supports up to 24 NVMe SSDs or 22 NVMe SSDs plus 2 SAS/SATA SSDs. It also supports 8 DDR4 DIMM slots and 128 lanes PCIe 3.0 offering the performance and density a variety of workloads.

ABOUT SAMSUNG NVME SSDS

[Samsung PM963 NVMe SSD](#), which is based on Samsung TLC Vertical NAND flash memory, is optimized for datacentre environments and engineered for outstanding performance under varying workloads. Leveraging the highly efficient NVMe, PM963 uses four PCIe lanes to deliver approximately four times the performance of SATA SSDs for sequential and random read, enabling significantly better performance-to-power efficiency compared to SATA SSDs.

ABOUT HPE M-SERIES

The 100GbE network includes [HPE M-series Ethernet switch](#) that is based on [Mellanox Spectrum™](#) switching controller, [ConnectX®-5 NICs](#) and [LinkX](#) cables and enables high performance workloads over Microsoft Windows Server 2016 Storage Spaces Direct. Using an end-to-end Mellanox 100GbE network solution enables the compute and storage traffic to run over a single high-performance network.

SUMMARY

The joint collaboration among HPE, Samsung, and Mellanox was able to successfully highlight the efficiency that is possible in Microsoft Windows Server 2016 based deployment using Storage Spaces Direct.

About Mellanox

Mellanox Technologies is a leading supplier of end-to-end InfiniBand and Ethernet interconnect solutions and services for servers and storage. Mellanox interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance capability. Mellanox offers a choice of fast interconnect products: adapters, switches, software, cables 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 and financial services.

To find out more, visit our website: www.mellanox.com



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