Migrating to Microsoft Storage Spaces Direct

King County (WA) Library System upgrades to Storage Spaces Direct and Hyper-V

Overview

King County Library System (KCLS), is one of the top library systems in the United States, serving over 1.4M residents in the region surrounding Seattle in Washington State. Recently King County Library System upgraded to Storage Spaces Direct (S2D) with DataON S2D-3110 solution. The S2D-3100 is a Microsoft-certified, all-flash NVMe SSD hyper-converged infrastructure, using Chelsio iWARP RDMA-enabled 10Gb Ethernet adapters as a low latency interconnect for the S2D cluster.

Microsoft Storage Spaces Direct is a storage feature introduced in Windows Server 2016 which enables building highly available and scalable storage systems by pooling local server storage. It allows you to build HA storage systems using networked storage nodes with only local storage, which can be disk devices that are internal to each storage node. It leverages SMB3 for all intra-node communication, including SMB Direct and SMB Multichannel, for low latency and high throughput storage. The Hyper-V VMs are then run on the same hosts as the storage in a hyper-converged deployment scenario.

Learn more about Storage Spaces Direct

DataON is the industry-leading provider of hyper-converged infrastructure and storage systems optimized for Microsoft Windows Server environments. The DataON S2D-3110 hyper-converged infrastructure, used in this case, is built to optimize the full stack of Microsoft Storage Spaces Direct in a hyper-converged platform. This appliance is built with integrated compute, network and storage infrastructure with near-linear scalability to simplify and maximize the deployment of Microsoft applications, virtualization, data protection and hybrid cloud services and provide the ultimate platform for the Microsoft software-defined data center (SDDC) solutions.

Learn more about the DataON S2D family of hyper-converged infrastructure

Chelsio 1/10/25/40/50/100Gb Ethernet Unified Wire adapters provide the optimal and scalable networking building block for Microsoft Storage Spaces Direct. Designed for industry-leading performance and efficiency, and with the unique ability to fully offload TCP/IP, iSCSI and iWARP protocols using a single ASIC and firmware, Chelsio adapters unburden communication responsibilities and processing overhead from servers and storage systems by enabling a true converged Unified Wire solution, resulting in a dramatic increase in application performance, and eliminates the need for a DCB enabled Ethernet switch to implement Storage Spaces Direct.

T520-CR iWARP RDMA enabled Unified Wire adapters, used in this case, also support the network QoS capability. Network QoS can be used in the hyper-converged configuration to ensure that the software-defined-storage system has enough bandwidth for communication between the nodes to ensure resiliency and performance.

One of the decisions KCLS needed to make was whether to use RDMA over Converged Ethernet (RoCE) or iWARP RDMA, as both RDMA networking protocols are supported by Windows Server 2016. KCLS chose to use Chelsio iWARP RDMA configuration because it does not require a lossless Ethernet network, and they were able to connect to their Cisco Nexus switch with no issues.

Storage Spaces Direct, Storage IOPS Performance with iWARP

Cluster configuration and setup details:

- **6 nodes of DataON S2D-3110 HCI**
  - 2 Intel® Xeon® E5-2600 v4 family processors in each node
  - 6 Chelsio T520-CR dual port 10Gb Ethernet Unified Wire adapters, one per node
  - 8 800GB 2.5" Intel DC P3700 NVMe SSDs
  - 32 2.0TB 2.5" Intel DC P3520 NVMe SSDs
- **Windows Server 2016 Storage Spaces Direct**
  - Cache: Intel DC P3700 NVMe SSDs
  - Capacity: Intel DC P3520 NVMe SSDs

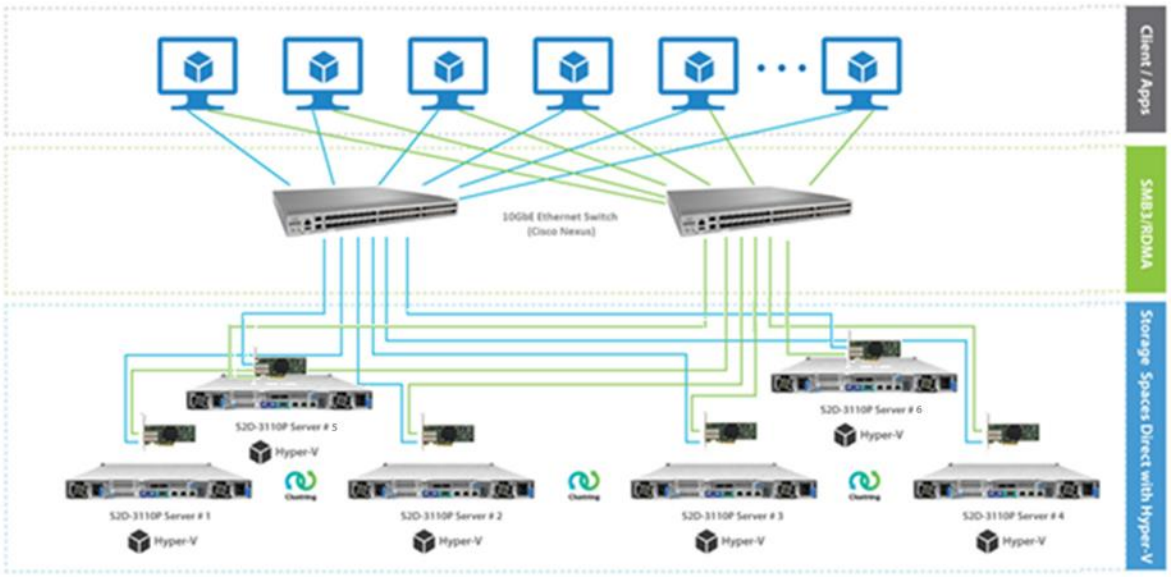

Figure 1 – Storage Spaces Direct Network Topology

Results

KCLS was able to consolidate three Hyper-V clusters into one, with the following benefits:

- 30% Increase in performance and now delivering 2.1M IOPS.
- Steady and predictable performance.
- Consolidated three Hyper-V clusters into one.
- Downsized their data center by 15-20 servers.

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	2,147,762	2,147,670	92	8,794	8,793	1		
hyperv17-1	383,953	383,937	16	1,570	1,570		0.206	0.647
hyperv17-2	369,471	369,468	2	1,513	1,513		0.233	0.916
hyperv17-3	412,819	412,789	30	1,691	1,691		0.240	0.715
hyperv17-4	536,600	536,582	18	2,198	2,198		0.486	1.296
hyperv17-5	444,918	444,894	24	1,822	1,822		0.410	0.825

Figure 2 – Storage Spaces Direct IOPS Numbers

In addition to supporting Storage Spaces Direct, the iWARP protocol also powers other aspects of Microsoft Windows such as **Storage Replica** for disaster recovery, **SMB Direct** for high performance file access, **Client RDMA** for bringing RDMA benefits to Windows 10 deployments, and **Network Direct** for Windows HPC deployments.

Summary

Chelsio iWARP RDMA-enabled 10Gb Ethernet adapter delivered a high-performance Storage Spaces Direct solution using standard Ethernet infrastructure and enabled customer to deploy Storage Spaces Direct with the appropriate high-performance network fabric. The ability to work with any non-DCBX switch, enabled an immediate plug and play deployment.

Support of iWARP protocol is enabled since Windows Server 2012R2 release, and is included in Windows Server 2016. This has allowed for years of testing for an extremely robust, tested and efficient deployment with Chelsio iWARP enabled Ethernet adapters. When you are using network as the transport for sharing server chassis, you don't want to cut corners or overcomplicate the solution. Chelsio's iWARP-based RDMA implementation is the only answer.

Related Links

Hyper-Converged Scale-Unit with Chelsio 40GbE
Storage Spaces Direct Performance with Network QoS
Storage Spaces Direct Performance with iWARP RDMA
iWARP RDMA – Best Fit for Storage Spaces Direct
High Performance Storage Spaces Direct with Chelsio 100GbE
Storage IOPS Update with Storage Spaces Direct - Microsoft Blog
Windows Server 2016 Storage Spaces Direct
Configuring Storage Spaces Direct