

## Hyper-Converged Scale-Unit with Chelsio 40GbE Chelsio T5 iWARP RDMA solution for Windows S2D + Hyper-V

### Overview

Microsoft **Storage Spaces Direct** (S2D) is a feature introduced in Windows Server 2016, which enables building highly available and scalable storage systems by pooling local server storage. You can now build HA storage systems using nodes with only local storage, which can be disks that are internal to each storage node. This not only eliminates the need for a shared SAS fabric and its complexities, but also enables using devices such as SATA solid state drives, which can help further reduce cost or NVMe solid state devices to improve performance. Storage Spaces Direct 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.

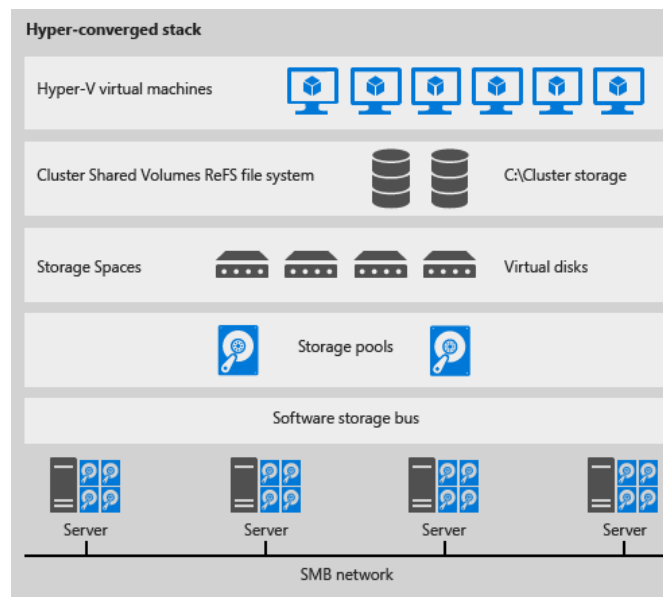


Figure 1 - S2D Hyper-Converged Stack

### Why Chelsio iWARP RDMA Solution for S2D

Chelsio's fifth generation, high performance RDMA 10/40Gbps Ethernet adapters, utilizing iWARP:

- Enable incremental, non-disruptive server installs.
  - Support the ability to work with any legacy (non-DCB) switch infrastructure.
  - Enable a decoupled server and switch upgrade cycle and a brownfield strategy to enable high performance, low cost S2D enablement.
- Are easy to use and install.

- Are cheaper to deploy → end user can purchase more compute servers for the same investment amount.
  - Do not require gateways or routers to connect to the TCP/IP world.
  - Save significant CPU cycles.
    - Enable cheaper CPUs for equivalent performance.
    - Enable significantly lower datacenter and utilities.
- Utilize very robust and stable protocols.
  - iWARP has been an IETF standard (RFC 5040) for 9 years, TCP/IP has been an IETF standard (RFC 793, 791) for 35 years.
    - No surprises, no fine print, plug and play.
  - Have multi-vendor support.
- Are supported in other Windows products.
  - Client RDMA in Windows 10 enables more deployment options.
  - Storage Replica is natively supported to enable disaster recovery.
  - Network Direct and Nano Server are natively supported.
  - iSCSI HW Offloaded initiator is natively supported.
- Are scalable to wherever the datacenter can scale to.
  - Inherit the loss resilience and congestion management from underlying TCP/IP.
- Are very high performance.
  - Extremely low latency, high bandwidth, high message rate.

## Storage Spaces Direct Performance Demonstration

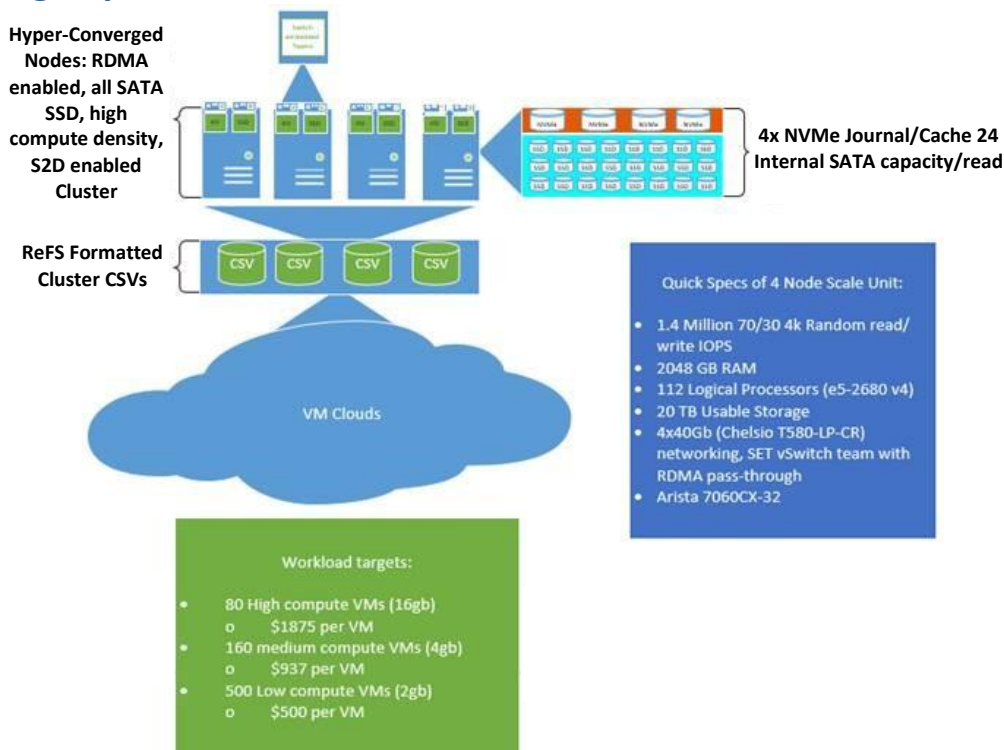


Figure 2 – S2D Performance Demonstration

This demonstration was engineered by ClearPointe and Raid Inc. using best of breed server hardware from Lenovo and HGST. It employs 4 nodes with SATA SSD and NVMe SSD connected across an iWARP RDMA-enabled fabric, using a 32 port Arista 7060CX switch. Each node was equipped with the following hardware:

- 2x Intel Xeon E5-2680v4 2.3Ghz (28c56t)
- 512GB RDIMM
- 4x 1.6TB HGST SN150 NVMe SSD (PCIe 3.0 x4)
- 3x LSI 9300 8i
- 24x 250GB Samsung DC Series evo SATA SSD
- 2x Chelsio T580-CR (Dual Port 40Gb PCIe 3.0 x8)
  - Chelsio Unified Wire driver v5.0.0.66
  - Chelsio FW v1.16.3.0
  - Dual port connected / adapter

A 40 virtual machines per node, for a total of 160 virtual machines configuration was used for this setup. Each virtual machine was configured with 2 vCPU and 4GB RAM. [VMFleet](#) tool was used to run [DISKSPD](#) in each of the virtual machines with various multi-threaded workloads from 4KiB to 32KiB IO size, mixed between 100 % READ and 100% WRITE, all random patterned.



**Figure 3 - Storage Spaces Direct Performance Numbers**

As seen from the above screenshot, this setup was able to demonstrate over 1.6M IOPS in aggregate using virtual machines, which means 10,000 IOPs per virtual machine!

***“This system represents the kind of well thought, best of breed, real-world high-performance solutions we provide our customers.”*** Daniel Weissenborn, Enterprise Architect, ClearPointe.



## Summary

**Chelsio RDMA enabled 40Gb Ethernet adapters** deliver a high performance Storage Spaces Direct (S2D) solution using standard Ethernet infrastructure and enables datacenters to deploy S2D with the right kind of high-performance network fabric. 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

[Storage IOPS Update with S2D - Microsoft Blog](#)

[5M IOPS with Storage Spaces Direct - Microsoft Tweet](#)

[ClearPointe case study using Chelsio iWARP RDMA Adapters](#)

[Windows Server 2016 Storage Spaces Direct](#)

[Configuring Storage Spaces Direct](#)