

Axellio demos WSSD cluster with Chelsio 100GbE

Great performance with Ease of Use

Introduction

Axellio demonstrated a four node WSSD cluster on Windows Server 2019 using Storage Spaces Direct (S2D) technology and Chelsio 100Gb iWARP RDMA adapters on their FabricXpress Platform. With a Chelsio T62100-CR adapter (2-port 100G RNIC) in each node and 24 Samsung NVMe devices per node, the cluster achieved 4.3 Million IOPS.

Results

Using vmfleet, 4K 100% sequential read was run and the following IOPs was observed.

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	4,322,354	4,322,354		17,702	17,702			
axnssd-node1	1,087,487	1,087,487		4,454	4,454		2.009	0.000
axnssd-node2	1,170,893	1,170,893		4,796	4,796		1.563	0.000
axnssd-node3	916,542	916,542		3,754	3,754		1.744	0.000
axnssd-node4	1,147,433	1,147,433		4,697	4,697		0.376	0.000

Figure 1 – S2D IOPS result

As can be seen, 4.3 Million IOPS was achieved!

Configuration

The demonstration was with a 4-node cluster as shown in the Figure 2. Each node had dual 22 core CPUs, 512GB of RAM per node, 24 x 800GB Samsung PM1725 NVMe drives and 1 T62100-CR adapter. 176 Virtual Machines (VMs) were created (44 per host), with 1 CPU core and 4GB of memory per VM.




Figure 2 – Topology of the Axellio FabricXpress HCI demonstration with Chelsio T62100-CR adapters

FX- WSSD starts at 4 nodes and scales up to 16 nodes per cluster.

FX-WSSD Appliance Specifications		
Software	Microsoft Windows Server 2019 Datacenter Edition	
Processors	Intel® XEON® E5-26xx V4 family	
Network	Chelsio10/25/40/50/100GbE	
Configurations	4 Node	16 Node
Memory	512GB to 4TB RAM	2TB to 16TB
Scalable NVMe Storage*	12.8TB to 920TB	51.2TB to 3.7PB
Form Factor	4U	16U

* Increased capacity supported as higher capacity SSDs become available

Chelsio iWARP RDMA does SMBDirect/S2D/WSSD

Chelsio’s T5 & T6 ASICs both have iWARP, Microsoft’s [preferred](#) flavor of RDMA, built-in. That means any Ethernet adapter Chelsio markets can efficiently run SMBDirect. These adapters enable high virtualized application performance by giving direct access to the data residing on hyper-converged or disaggregated S2D storage, as well as higher scalability for S2D. Chelsio iWARP also enables much more streamlined and simplified configuration of RDMA high-performance networking for Microsoft S2D compared to other Ethernet RDMA networking alternatives. Microsoft also recommends iWARP as the safer alternative for S2D as it does not require any configuration of DCB on network hosts or network switches and can operate over the same distances as any other TCP connection.

The available adapters from Chelsio are based on these two ASICs and depending on the model, can achieve 1/10/25/40/50/100 Gb Ethernet speeds. In this demonstration, Axellio’s FabricXpress HCI demonstrated with one of several 100Gb models, the T62100-CR.




Figure 3 – Chelsio’s T62100-CR

Conclusion

Chelsio's adapters are perfect for environments that demand high performance. Using the T62100-CR adapters, Axellio FabricXpress WSSD HCI has achieved an outstanding 4.3 Million IOPS using standard Ethernet infrastructure. The ability to work with any non-DCBX switch, enables an immediate plug and play deployment.

Related Links

[WSSD Configuration using iWARP Mesh Topology](#)

[Axellio demos WSSD cluster at Microsoft Ignite](#)

[Storage Spaces Direct throughput with 100GbE iWARP - Microsoft Blog](#)

[Migrating to Microsoft Storage Spaces Direct](#)

[S2D Performance with Network QoS](#)

[iWARP RDMA – Best Fit for Storage Spaces Direct](#)

[High Performance S2D with Chelsio 100GbE](#)