Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

The Marvell® FastLinQ® 41000 Series is a high-performance 25GbE adapter with iSCSI hardware offload and Universal RDMA performance benefits.

Executive Summary

Non-Volatile Memory Express (NVMe) and Storage Class Memory (SCM) are offered in current generation servers with Intel Xeon Scalable Processors. The resulting server storage performance gains are driving an increase in network bandwidth: Virtual Machines (VMs) and containers are more densely deployed on servers, the internet Small Computer System Interface (iSCSI) is being used for high-bandwidth storage solutions, and Hyper Converged Infrastructure (HCI) needs extensive bandwidth for inter-node communications. 25GbE networking is not only necessary, but the new standard.

25GbE has become a cost-effective upgrade to 10GbE as it delivers 2.5 times the data throughput of 10GbE through clock speed alone. Unlike 40GbE and others, extra lanes are not necessary, so 25GbE is backwards compatible with 10GbE.

In addition to speed, it is also important to consider additional Network Interface Card (NIC) features that enhance overall performance that may be applicable to a deployment. This includes features like iSCSI hardware offload or Remote Direct Memory Access (RDMA). Marvell FastLinQ 41000 Series 25GbE adapters support these and many other capabilities.

Marvell commissioned Demartek to evaluate the benefits of the Marvell FastLinQ 41000 Series when used with latest generation servers. We tested for Layer 2 performance, compared the iSCSI hardware initiator offload performance to that of software initiator on a leading competitor, evaluated Marvell FastLinQ 41000 Series use in a hyper-converged Storage Spaces Direct (S2D) cluster with SCM and NVMe storage, and weighed the benefits of upgrade by comparing 25GbE performance with 10GbE performance.

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Key Findings

- The Marvell FastLinQ 41000 Series achieved line rate bidirectional 25GbE performance for buffer sizes of 2KB up to 1MB, offering strong handling of peaks and scalability for enterprise applications.
- The Marvell FastLinQ 41000 Series hardware iSCSI initiator usually achieved line rate throughput on read and write for all block sizes 4KB and above while the Linux software initiator on the Mellanox adapter was only able to consistently achieve line rate on writes.
- In all Microsoft Windows test cases, the Marvell FastLinQ 41000 Series hardware iSCSI initiator consumed less processor than the Microsoft Windows software initiator on Mellanox, especially for larger block sizes.
- The 25GbE Marvell FastLinQ 41000 Series 25GbE adapter achieved on average 126% more read throughput and 62% more write throughput than the same cluster using 10GbE in a S2D test.
- For large block S2D read testing, the cluster utilizing the Marvell FastLinQ 41000 Series 25GbE adapter with Universal RDMA achieved a total average throughput of 23,714 MBPS while using on average 41% of available cluster processor.
- For large block S2D write testing, the cluster utilizing the Marvell FastLinQ 41000 Series 25GbE adapter with Universal RDMA achieved a total average throughput of 2,903 MBPS, while using on average 13% of available cluster processor
- The 25GbE technology consistently outperformed the 10GbE technology for throughput in iSCSI tests in Linux and Windows environments.
- Marvell's ability to support both RoCE and iWARP RDMA provide customers with flexibility on their networking implementation which can reduce implementation time and/or simplify network complexity and management.

Marvell FastLinQ 41000 Series 25GbE

The Marvell FastLinQ 41000 Series supports:

- High bandwidth 25GbE
- Backwards compatible with 10GbE to aid in 10GbE to 25GbE migration
- Marvell SmartAN™ technology automatically detects connected cables and switch to enable zero touch transition between 10G and 25G
- Single Root Input/Output Virtualization (SR-IOV) and Network Partitioning (NPAR)
- Full iSCSI hardware offload
- Universal RDMA support for all current protocols:
    - RDMA over Converged Ethernet (RoCE)
    - RoCE version 2 (RoCEv2)
    - Internet Wide-area RDMA Protocol (iWARP)
- Forward Error Correction (FEC) – RS-FEC and FC-FEC
- Tunnel Offload – NVGRE, VXLAN, GENEVE
- DPDK Small Packet Acceleration

Marvell adapters have excellent performance and support the widest arrays of features on the market.

In addition, the 25GbE Marvell FastLinQ 41000 Series adapter is compatible for future generations of servers with PCIe Gen4.

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Marvell FastLinQ 41000 Series Layer 2 Performance

Two Marvell FastLinQ 41000 Series adapters were installed in two current generation servers with Intel Xeon Scalable processors, formerly known as “Purley” platforms. One port from each server was connected to a switch and iperf was used to test bandwidth between the two servers. Buffer sizes from 512B up to 1MB were tested at varying thread counts.

Both bidirectional and unidirectional bandwidth was measured. In all but the smallest buffer sizes, we saw line rate 25GbE performance.

The Marvell FastLinQ 41000 Series adapter was then reconfigured to 10GbE and the test was run again. As expected, 25GbE had 2.5 times the performance of 10GbE. This translates to better handling of peaks and scalability for enterprise applications.

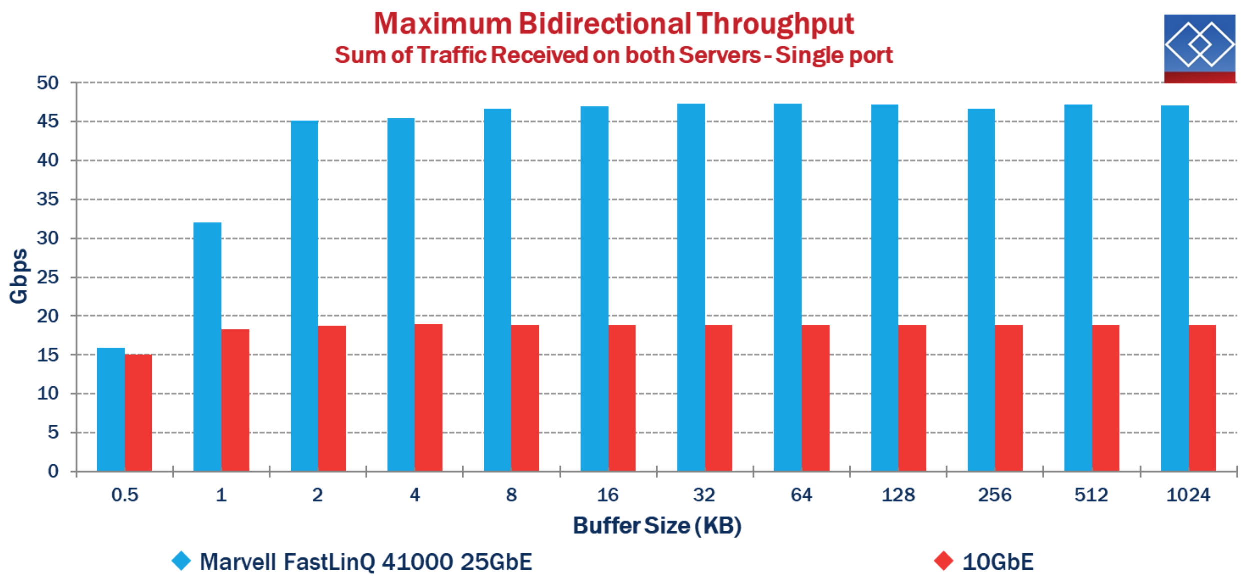

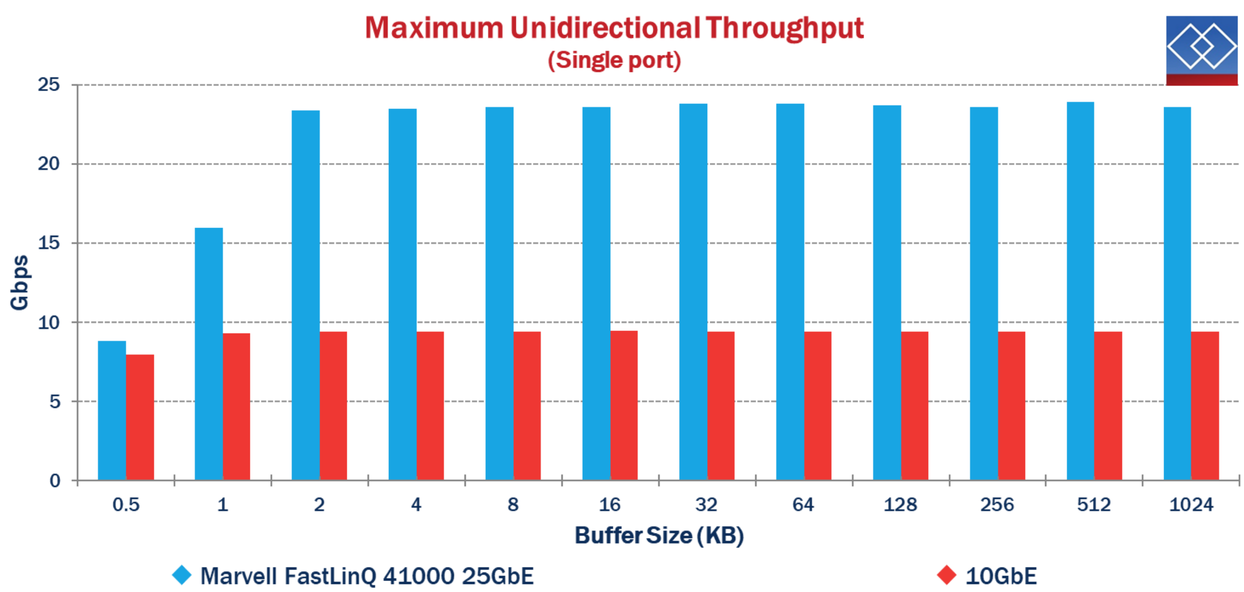

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Marvell FastLinQ 41000 Series iSCSI Hardware Offload

Software versus Hardware Initiators

The iSCSI initiator can be implemented in software or hardware. Software initiators are usually provided by the operating system and use system resources and any Network Interface Card (NIC) available. Hardware offloaded Initiators are available on select NICs. Hardware Initiators usually have their own TCP/IP stack and offload iSCSI processing to the NIC, conserving valuable system resources. Hardware Initiators are most valuable in deployments with limited processor capacity. Hardware initiators will reduce system processor utilization and, in some cases, can improve IOPS and throughput. The Mellanox ConnectX-4 LX does not offer hardware iSCSI initiator, however a fully hardware offloaded initiator is available on the Marvell FastLinQ 41000 Series adapter.

Performance Test Setup

Tests comparing the Marvell FastLinQ 41000 Series hardware iSCSI initiator to the operating system software iSCSI initiator with the Mellanox ConnectX-4 Lx were performed to validate the performance benefits offered by Marvell's hardware initiator. 32 iSCSI storage targets were deployed. A Marvell FastLinQ 41000 Series Adapter was deployed in the test server and the hardware initiator was used to connect to the iSCSI storage targets. Tests were run with fio on Linux and with Diskspd on Windows. The Marvell FastLinQ 41000 Series adapter was then configured to 10GbE and the test was run again. After this the Marvell Adapter was replaced with an Mellanox ConnectX-4 LX 25GbE and the tests were repeated a third time using the operating system provided software initiator.

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Marvell FastLinQ 41000 Series Hardware iSCSI Initiator versus Linux Software Initiator on Mellanox ConnectX-4 Lx

For 512B random reads, the Marvell FastLinQ 41000 Series hardware iSCSI initiator achieved 113% more topline IOPS than the Linux software initiator with Mellanox, and an average of 87% more IOPS for all unidirectional workloads.

The overall max IOPS achieved by the Marvell FastLinQ 41000 Series hardware iSCSI initiator was 82% higher than the overall max IOPS achieved by the Linux software initiator on Mellanox.

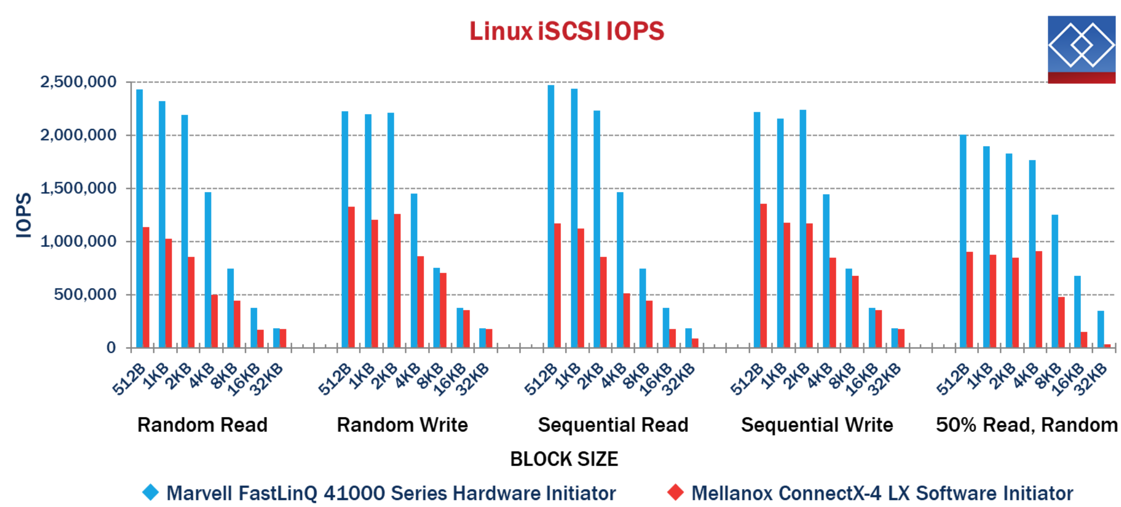

The Marvell FastLinQ 41000 Series hardware iSCSI initiator achieved IOPS that routinely outpaced the Linux software initiator on Mellanox by a wide margin. Despite this, in 80% of our tests, the Marvell hardware initiator utilized fewer processor cycles. This is typical of the processor savings usually seen in hardware offload. For example, the Marvell FastLinQ 41000 Series hardware iSCSI initiator achieved 67% more IOPS than Mellanox while using 40% less processor for 4KB random writes. Without hardware offload, we would expect the processor utilization to increase with IOPS, possibility limiting the scalability of enterprise applications.

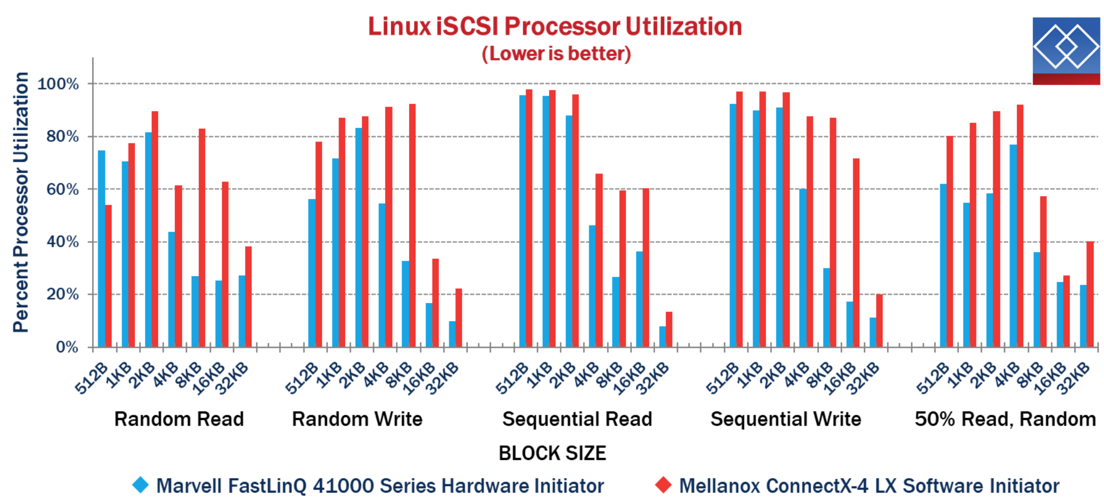

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Calculating processor effectiveness can also help us to see the processor savings from hardware offload. We define processor effectiveness as the ratio of IOPS or throughput to percent processor utilization. This effectively tells us for each 1% of processor utilized, how much work can be achieved:

> The Marvell FastLinQ 41000 Series hardware iSCSI initiator used 65% less processor resources than the Linux software initiator on Mellanox did to achieve line rate for 8KB random writes, achieving 3 times the processor effectiveness of Linux software initiator on Mellanox.

> The Marvell FastLinQ 41000 Series hardware iSCSI initiator achieved 193% more IOPS than Linux software initiator on Mellanox did while using 29% less processor for 4KB random reads, achieving 4.1 times the processor effectiveness of Linux software initiator on Mellanox.

> The Marvell FastLinQ 41000 Series hardware iSCSI initiator achieved 66% more IOPS than Linux software initiator on Mellanox did while using 67% less processor for 8KB random reads, achieving 5.1 times the processor effectiveness of Mellanox.

The Marvell FastLinQ 41000 Series iSCSI hardware initiator achieved line rate throughput on read and write for all but the smallest block sizes. Linux software initiator on Mellanox was only able to consistently achieve line rate on writes.

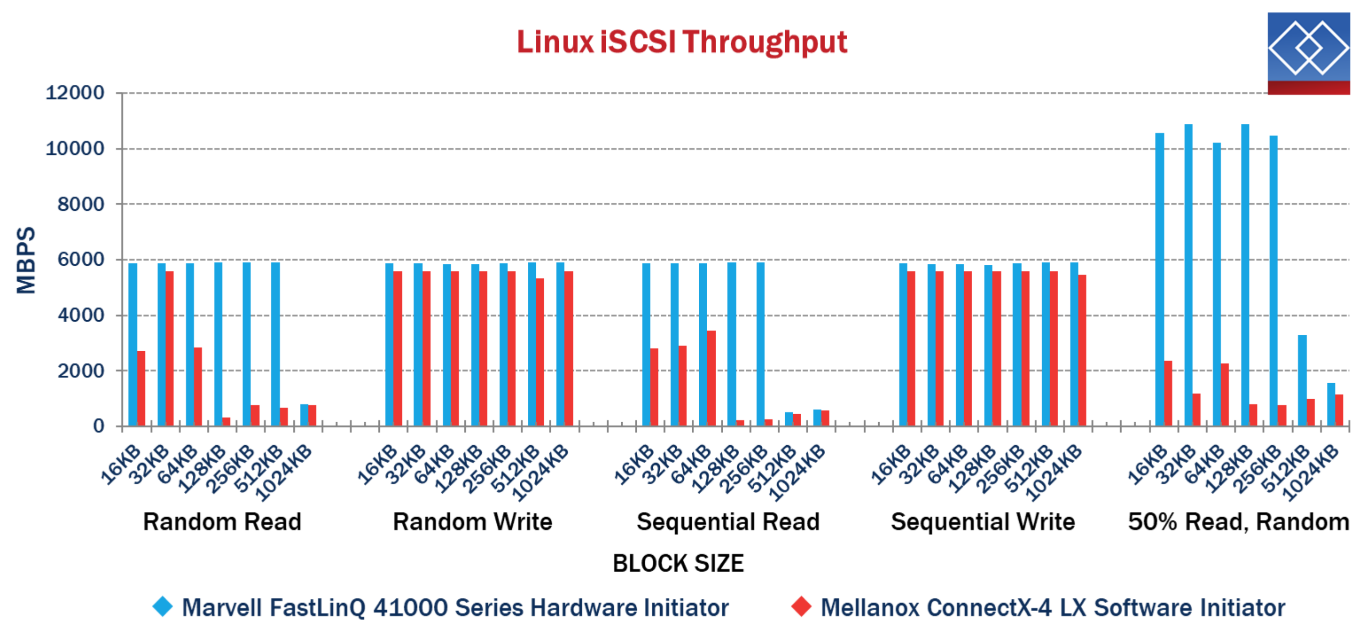

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Demartek
July 2019

Marvell FastLinQ 41000 Series Hardware iSCSI Initiator versus Windows Software Initiator on Mellanox ConnectX-4 Lx

For our smallest block sizes, the Marvell FastLinQ 41000 Series hardware iSCSI initiator achieved much higher IOPS with only small additional processor expense when compared to the Microsoft Windows software initiator on Mellanox. For example, for all 512B testing, the Marvell FastLinQ 41000 Series hardware initiator achieved on average 30% more IOPS while consuming on average 27% less processor.

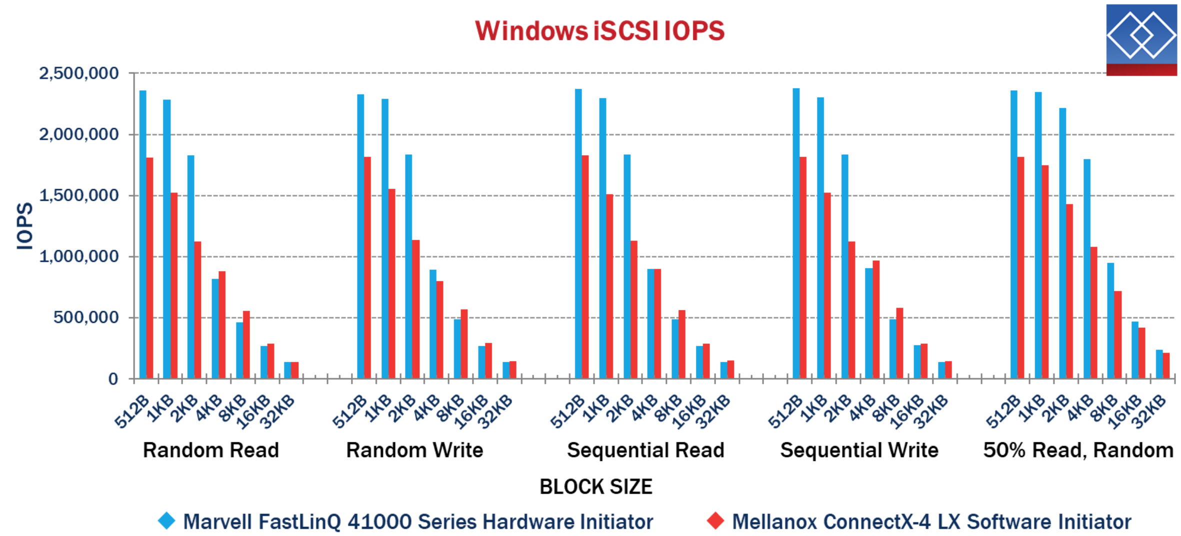

For 4KB 50% read, the Marvell FastLinQ 41000 Series hardware iSCSI initiator achieved 66% more IOPS while consuming 9% less processor than the Microsoft Windows software initiator on Mellanox.

Similarly, for 8KB 50% read, the Marvell FastLinQ 41000 Series hardware iSCSI initiator achieved 33% more IOPS while consuming 30% less processor than the Microsoft Windows software initiator on Mellanox.

In all test cases, the Marvell FastLinQ 41000 Series hardware iSCSI initiator consumed less processor than the Microsoft Windows software initiator on Mellanox.

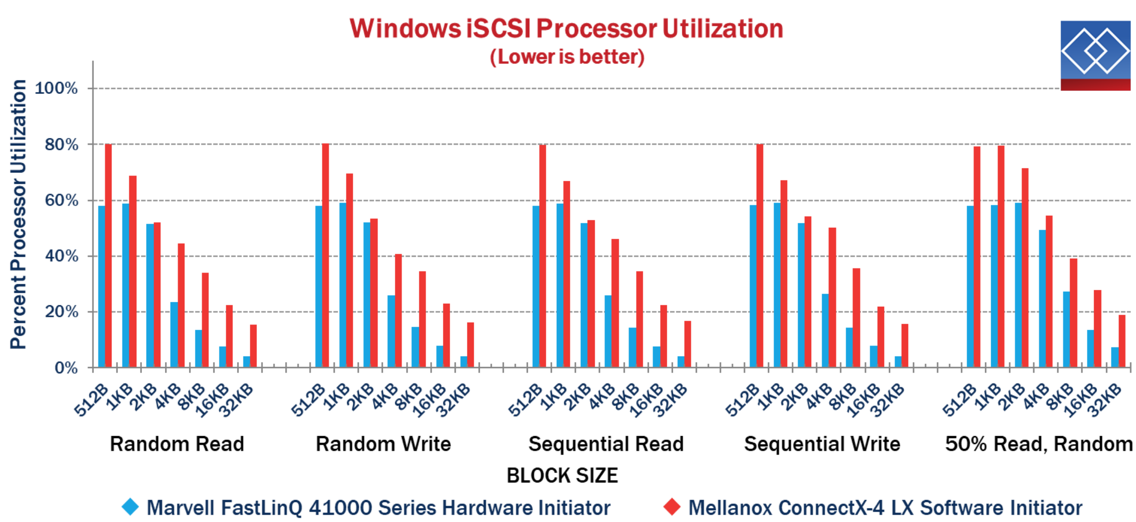

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

The Marvell FastLinQ 41000 Series hardware initiator was close to line rate throughput for all but the smallest block sizes.

It should be noted that the software initiator consumes considerably more CPU cycles to achieve its performance for the larger block sizes.

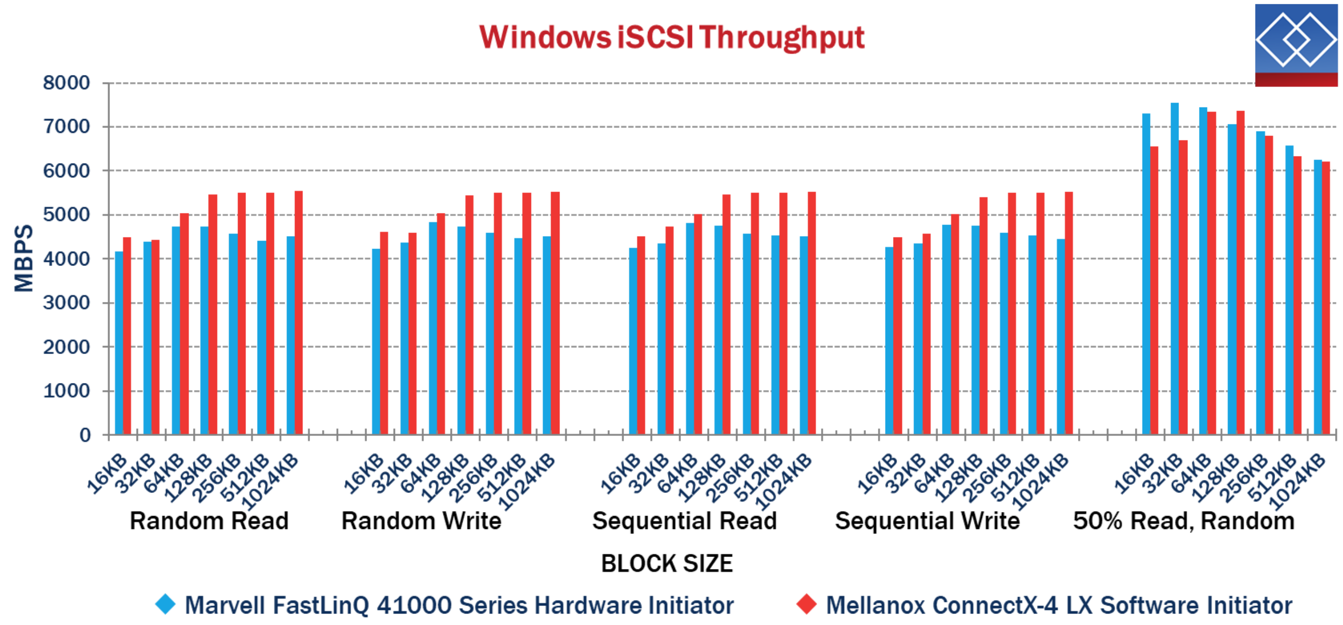

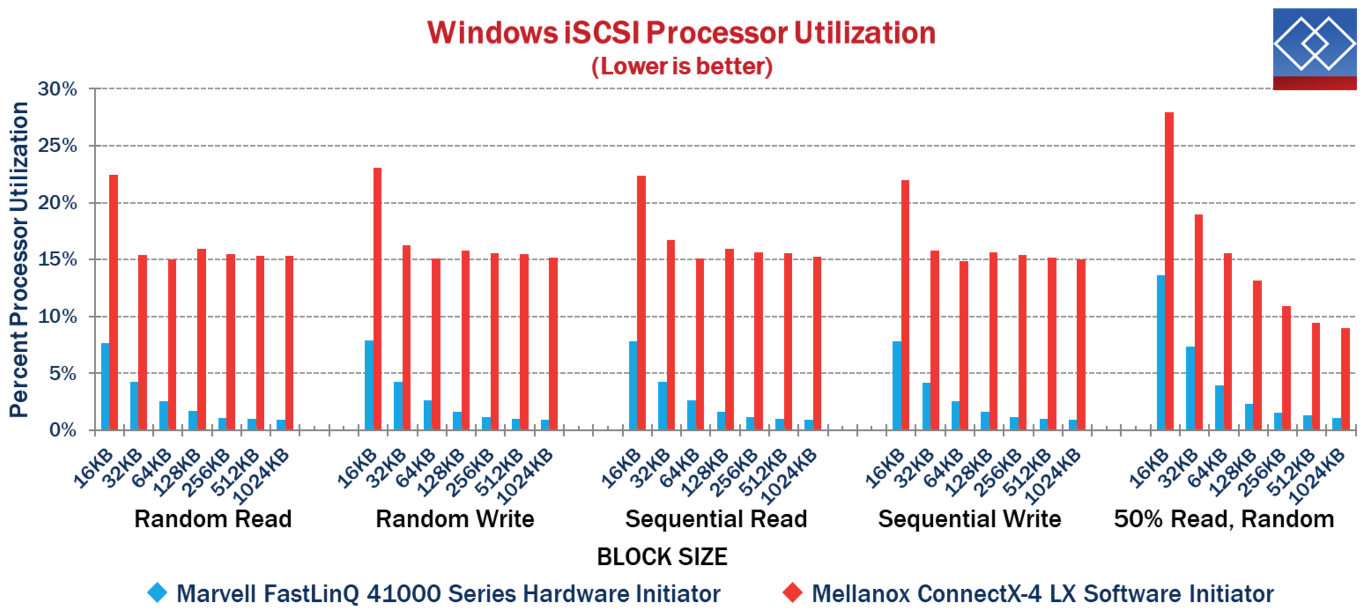

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Marvell FastLinQ 41000 Series Hardware iSCSI Initiator 25GbE Performance versus 10GbE Performance

The performance benefits of 25GbE over 10GbE translate to 2.5 times the bandwidth for iSCSI workloads.

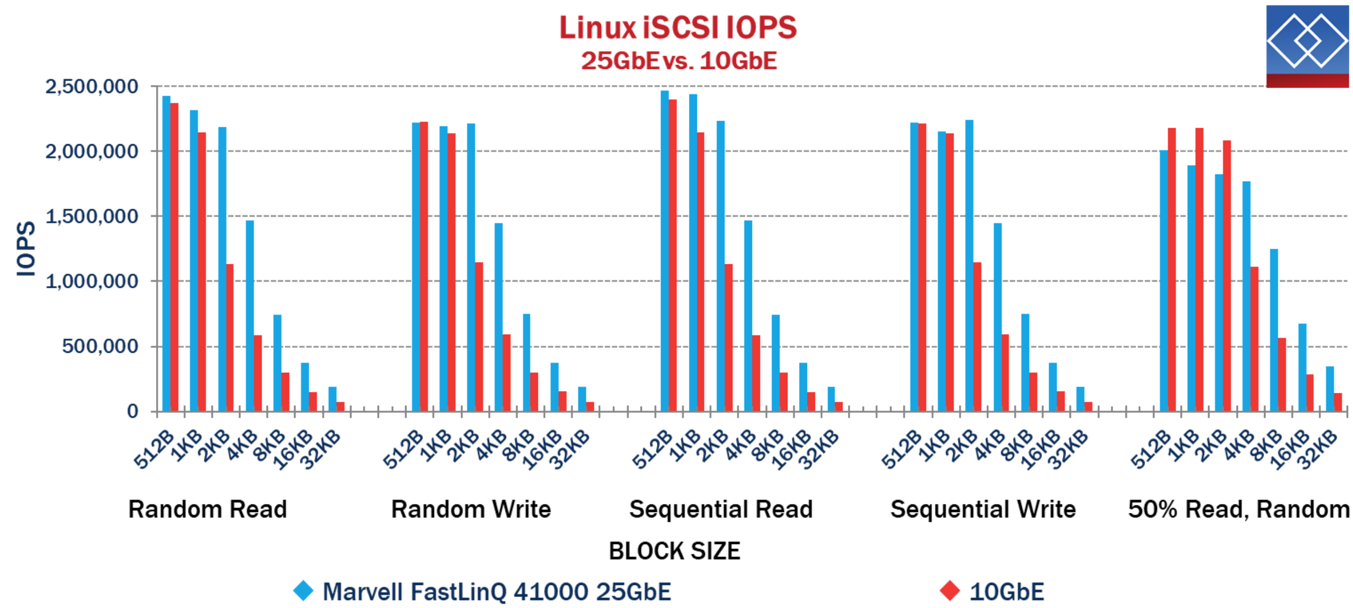

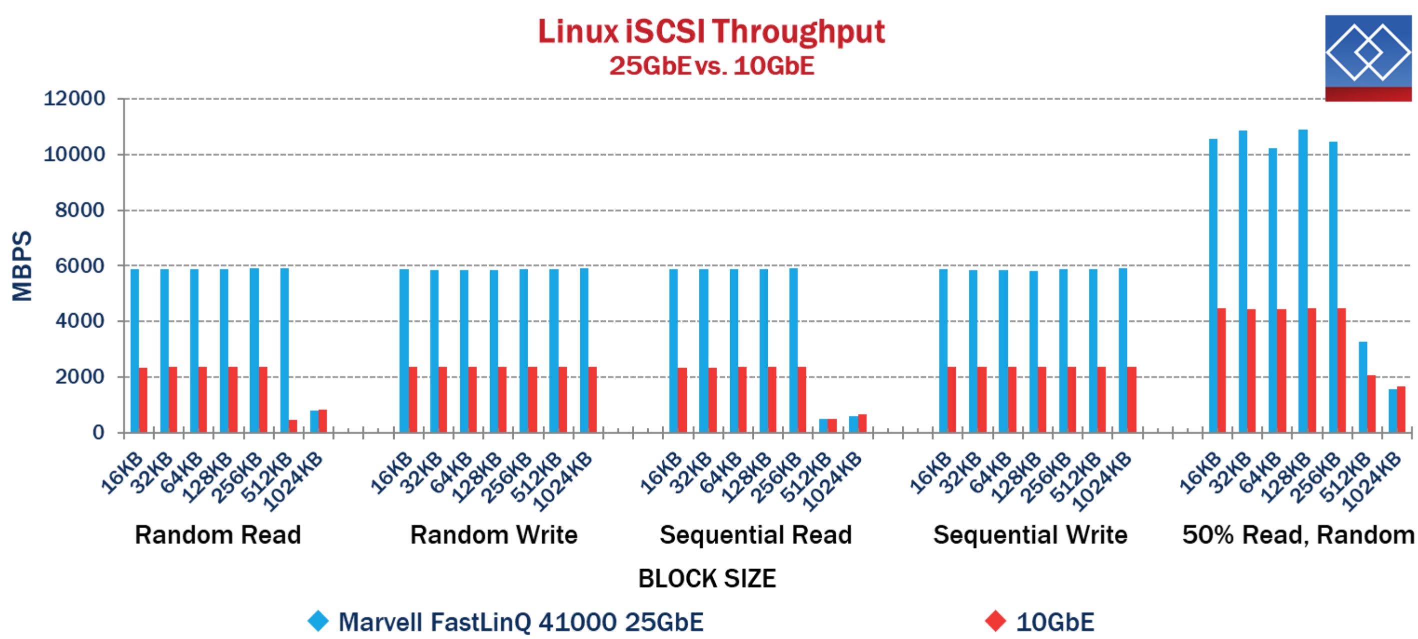

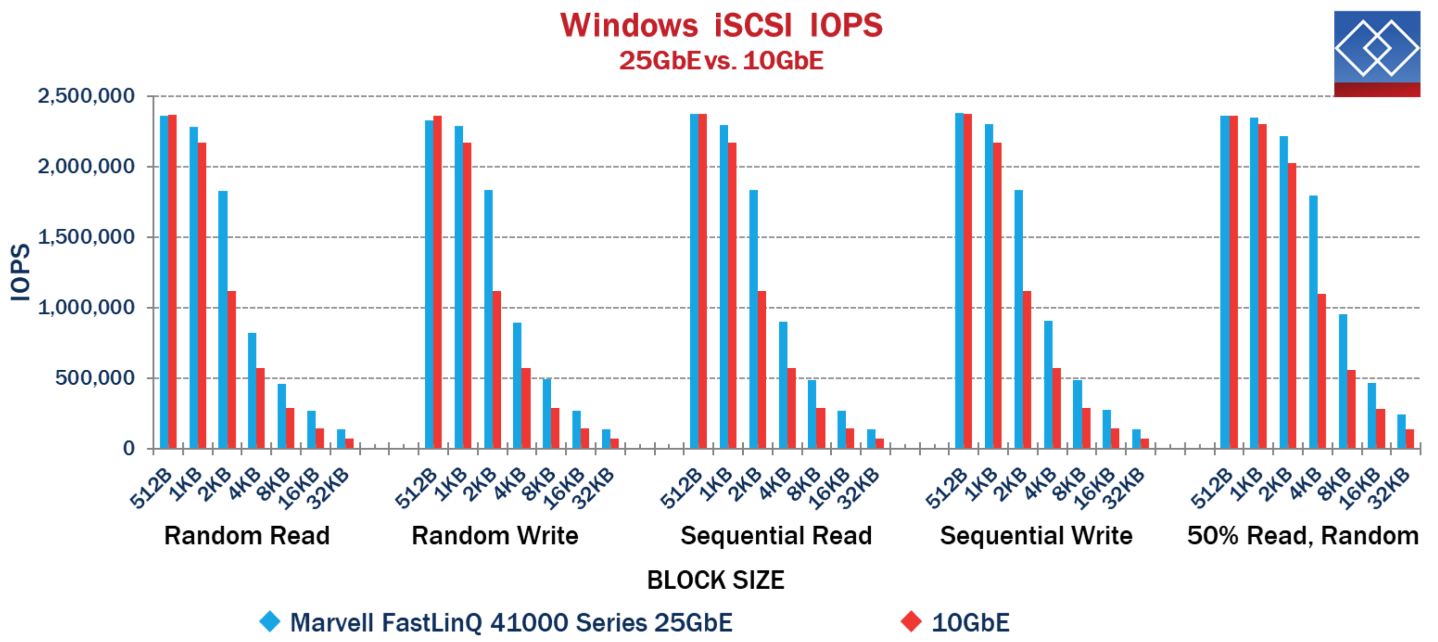
Windows iSCSI IOPS
25GbE vs. 10GbE
2,500,000
2,000,000
1,500,000
1,000,000
500,000
0
IOPS
512B 1KB 2KB 4KB 8KB 16KB 32KB
Random Read
Random Write
Sequential Read
Sequential Write
50% Read, Random
BLOCK SIZE
Marvell FastLinQ 41000 Series 25GbE
10GbE

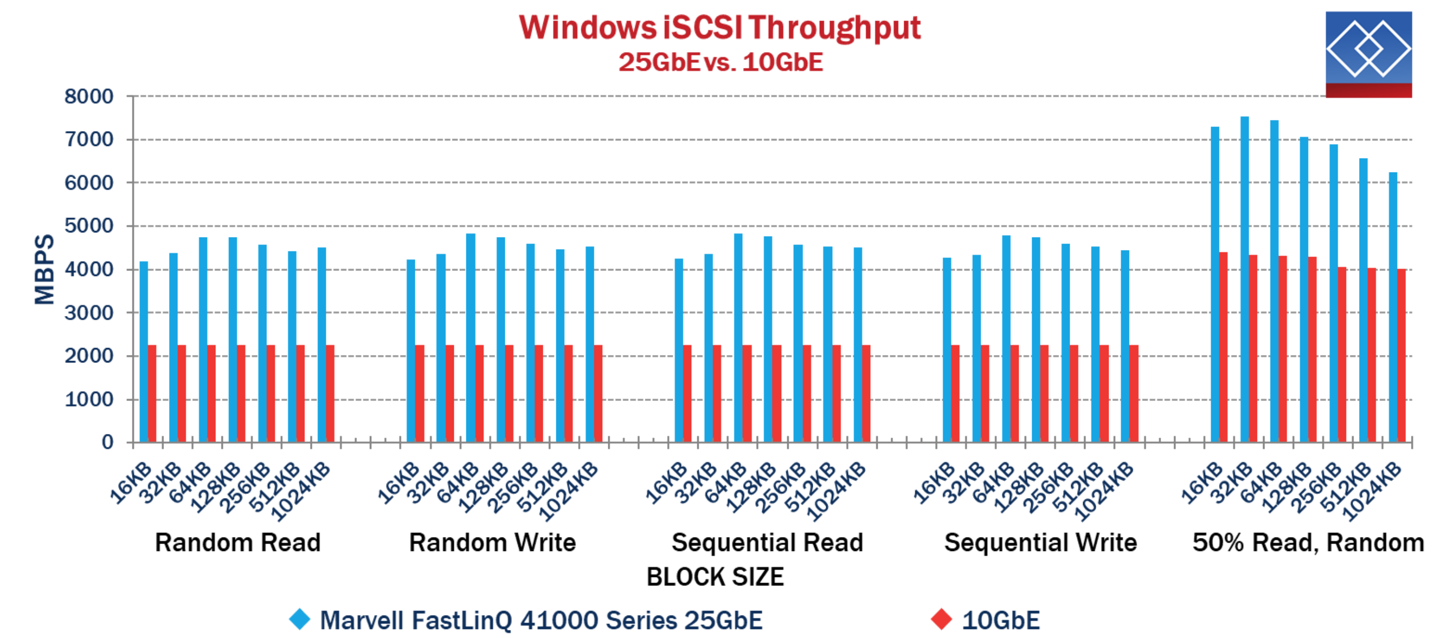
Windows iSCSI Throughput
25GbE vs. 10GbE
8000
7000
6000
5000
4000
3000
2000
1000
0
MBPS
16KB 32KB 64KB 128KB 256KB 512KB 1024KB
Random Read
Random Write
Sequential Read
Sequential Write
50% Read, Random
BLOCK SIZE
Marvell FastLinQ 41000 Series 25GbE
10GbE

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Marvell FastLinQ 41000 Series RDMA

RDMA is the remote memory management capability that allows server-to-server data movement directly between the application memory of each server without any processor involvement. Ethernet-based RDMA requires specialized NICs, sometimes called RNICs.

RoCE and iWARP

RDMA has two main implementations: RoCE and iWARP. For best results, RoCE deployments typically need a lossless fabric supplied by a switch that supports Data Center Bridging (DCB). A combination of Flow Control, Priority Flow Control (PFC), Enhanced Transmission Selection (ETS) and Explicit Congestion Notification (ECN) may be used to improve network performance and guarantee losslessness for RoCE traffic. Care must be taken to determine the optimal configuration for deployment and all switches and adapters must be configured identically as a single dropped packet can be extremely detrimental to RoCE performance. However, when deployed correctly RoCE can provide the most efficient performance.

iWARP does not require a lossless fabric, making deployment simpler as the standard TCP/IP stack is used. However, latency delivered by iWARP could be higher than RoCE, but most deployments like S2D would be unaffected due to this difference.

Both types of RDMA may be used in many use cases, including iSCSI Extensions for RDMA (iSER), Network File System over RDMA (NFSoRDMA), NVMe over Fabrics (NVMe-oF), and Server Message Block Direct (SMB Direct). Microsoft uses SMB Direct in an HCI environment in Windows Server 2016's and 2019's Storage Spaces Direct (S2D), where RDMA traffic is used for cluster inter-node communications. S2D is positioned as the top use case for RDMA – both RoCE and iWARP. Microsoft recommends iWARP for S2D (*).

Performance Test Setup

A four node Storage Spaces Direct Cluster was created. Each server node had the following hardware:

> 1xMarvell FastLinQ 41000 Series Adapter.

> 4xNVMe drives.

> 6x16GB NVDIMMs (used as SCM).

* https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/storage-spaces-direct-hardware-requirements

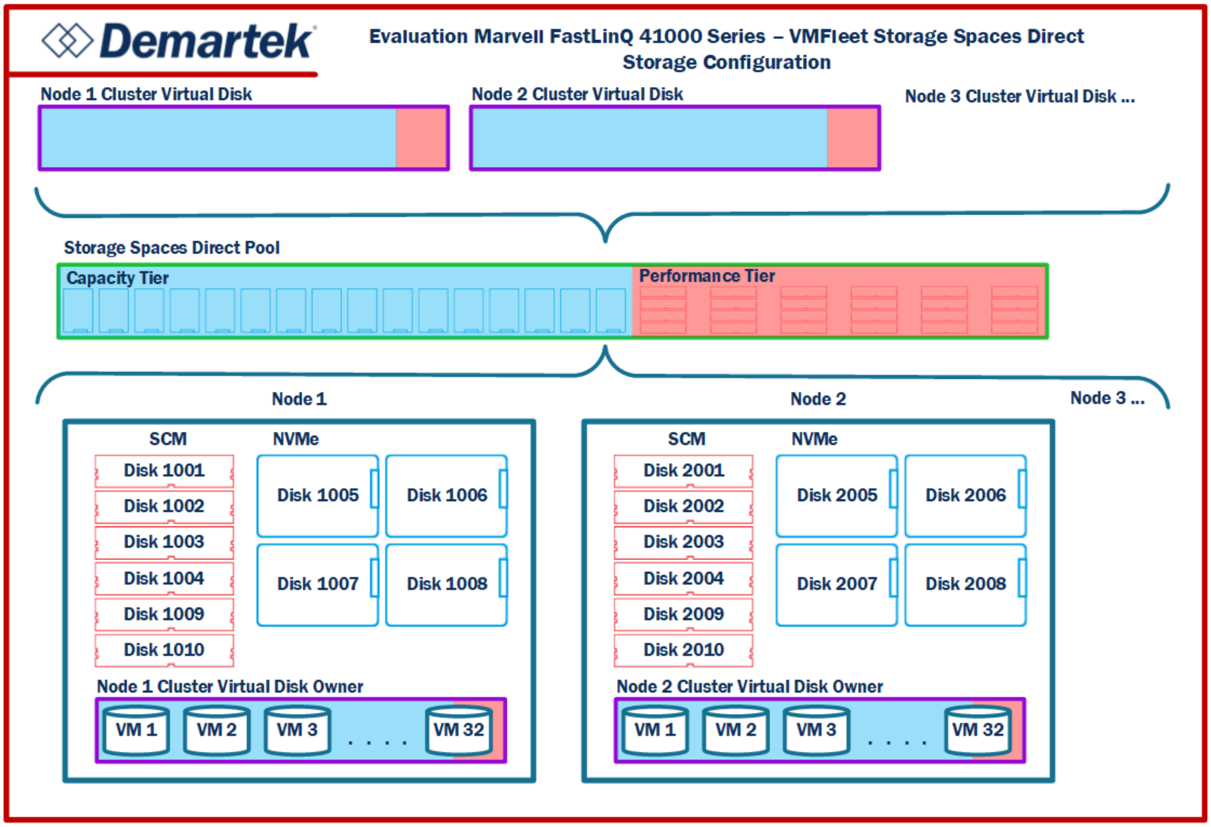

An S2D cluster with 3-way mirroring was created from these servers with the SCM NVDIMMs as the performance tier and NVMe as the capacity tier.

Switch Embedded Teaming (SET) with RDMA was used for inter-node communication.

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

July 2019

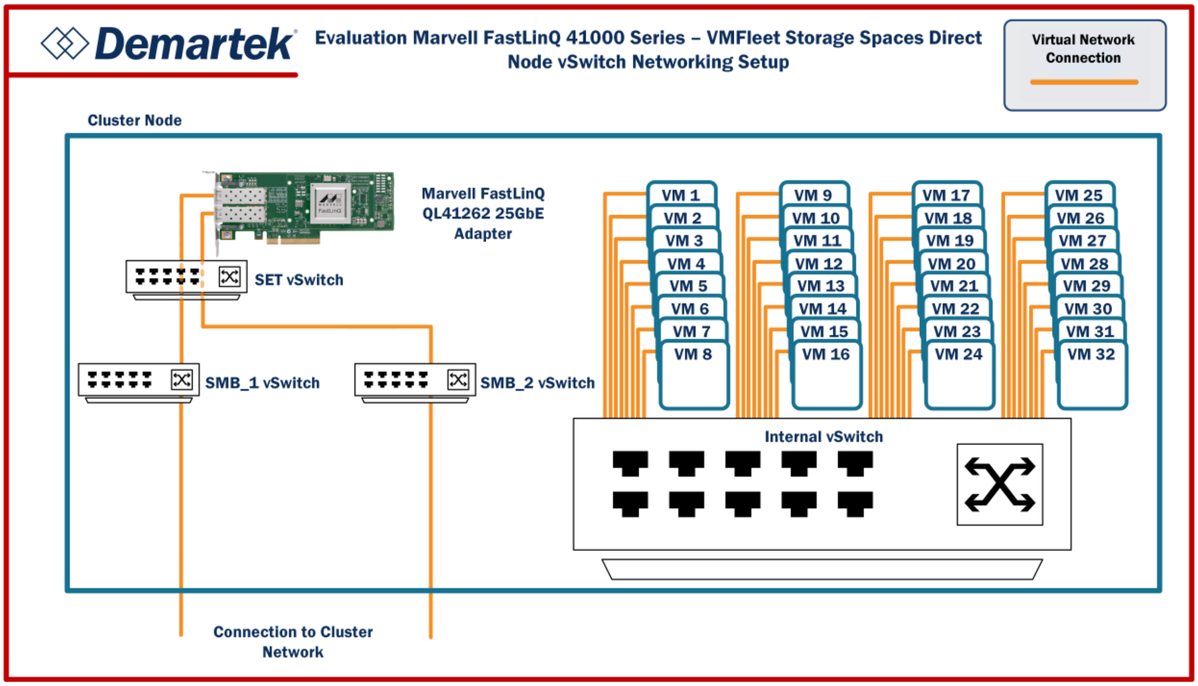

VMfleet is an open source tool developed by Microsoft and is used to load and stress servers, typically in order to characterize S2D performance. It consists of VMs deployed on the S2D Storage, and scripts to run coordinated DiskSpd test runs across the VMs, was deployed on the cluster and the total cluster throughput was measured.

VMfleet generates inter-node traffic that is typical of hyper-converged environments, producing high-performance on read and lower performance on write. As with any environment, the underlying storage media influences the behavior of the environment. SSDs typically have poorer write performance than read performance. However, in addition to this, S2D's 3-way mirroring will make additional work when a cluster writes. Upon write, three copies of the item must be written to three different nodes. This generates quite a bit of inter-node traffic, and greatly increases the work of each write.

By contrast, upon read, only one copy of the item must be read. In most of these cases, because there are so many copies of the data in the 3-way mirror, there is a high chance the data will be local, and limited network traffic is generated. Lastly, cluster performance will be affected by our S2D storage tier structure.

Workloads with both random and sequential access patterns were run, as is standard in most storage performance tests, however, the results showed very little performance difference between the access patterns. For this reason, only random performance is shown here.

The same set of tests were run four times. Once with RoCE on 25GbE, once with iWARP on 25GbE, once with RoCE on 10GbE, and once with iWARP on 10GbE.

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Read Performance

For small block reads, iWARP had a higher latency than RoCE. But when block size increased, RoCE had higher latency. However, the differences between them could be considered in the realm of noise.

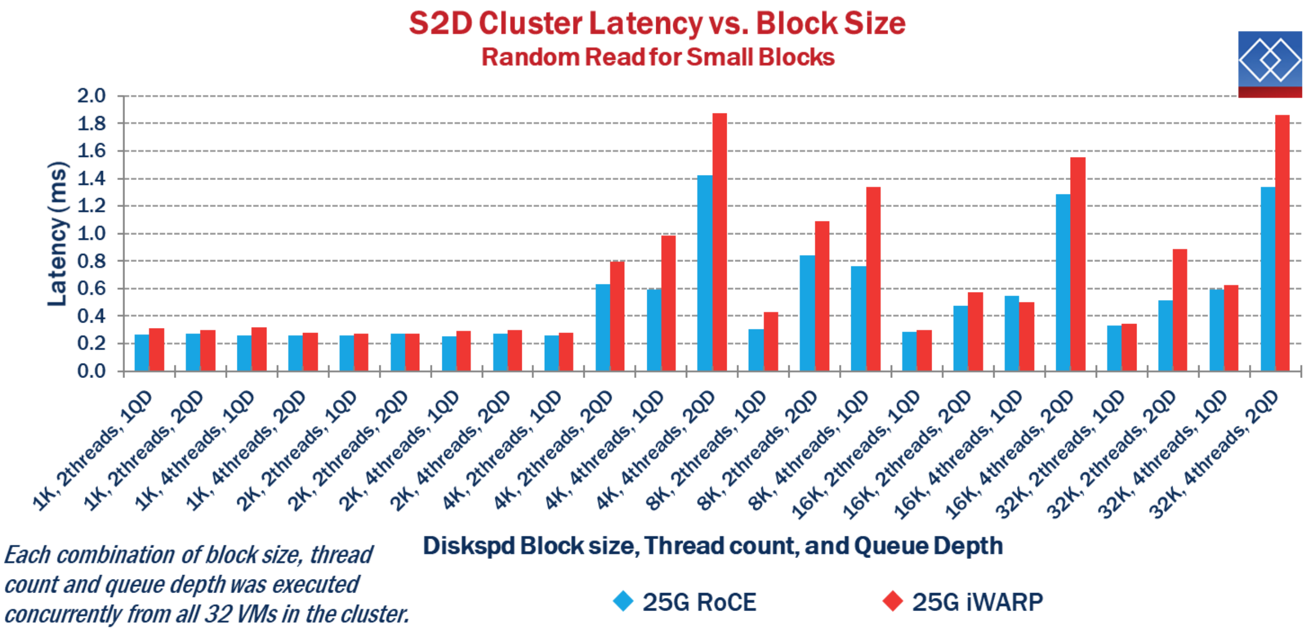

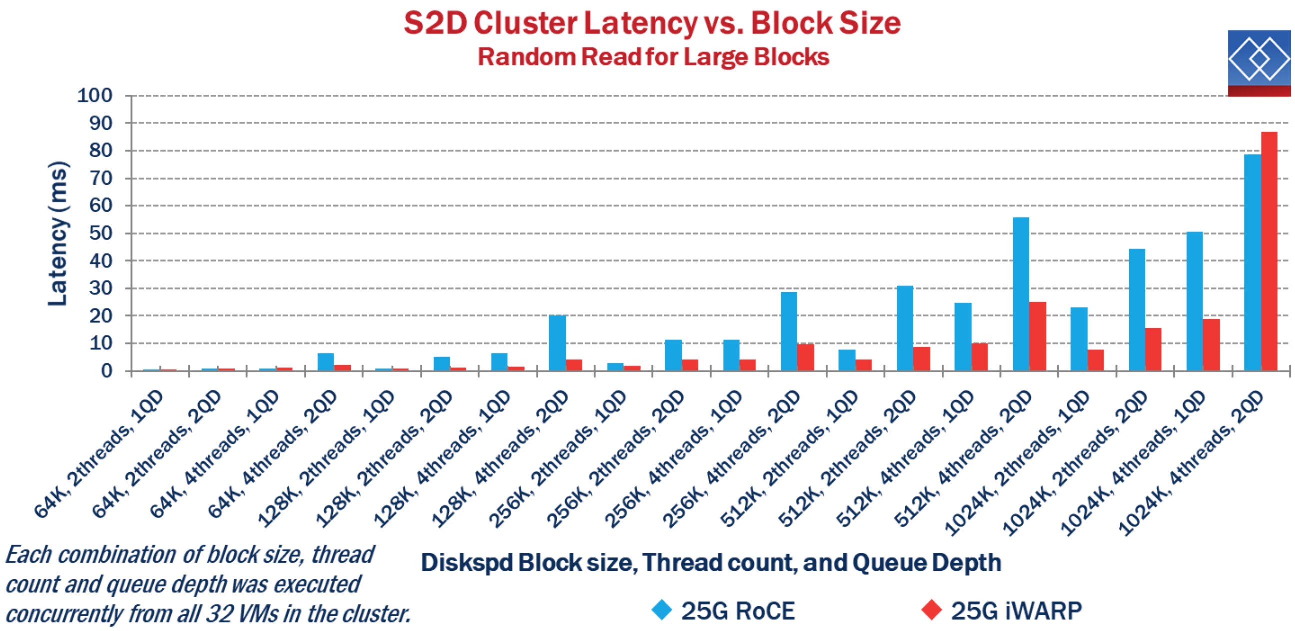

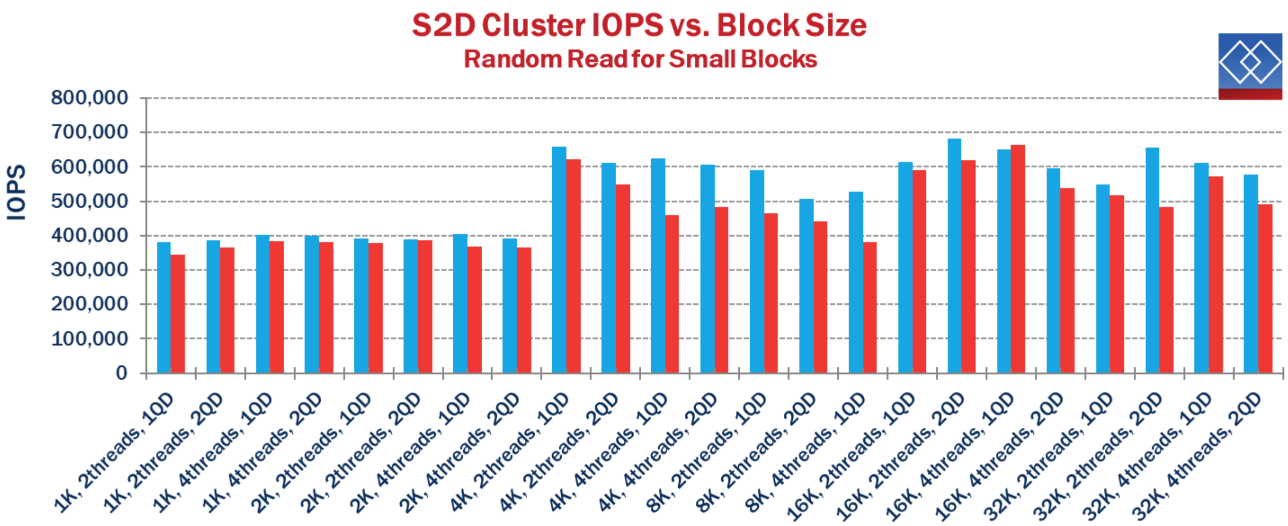

Each combination of block size, thread count and queue depth was executed concurrently from all 32 VMs in the cluster.

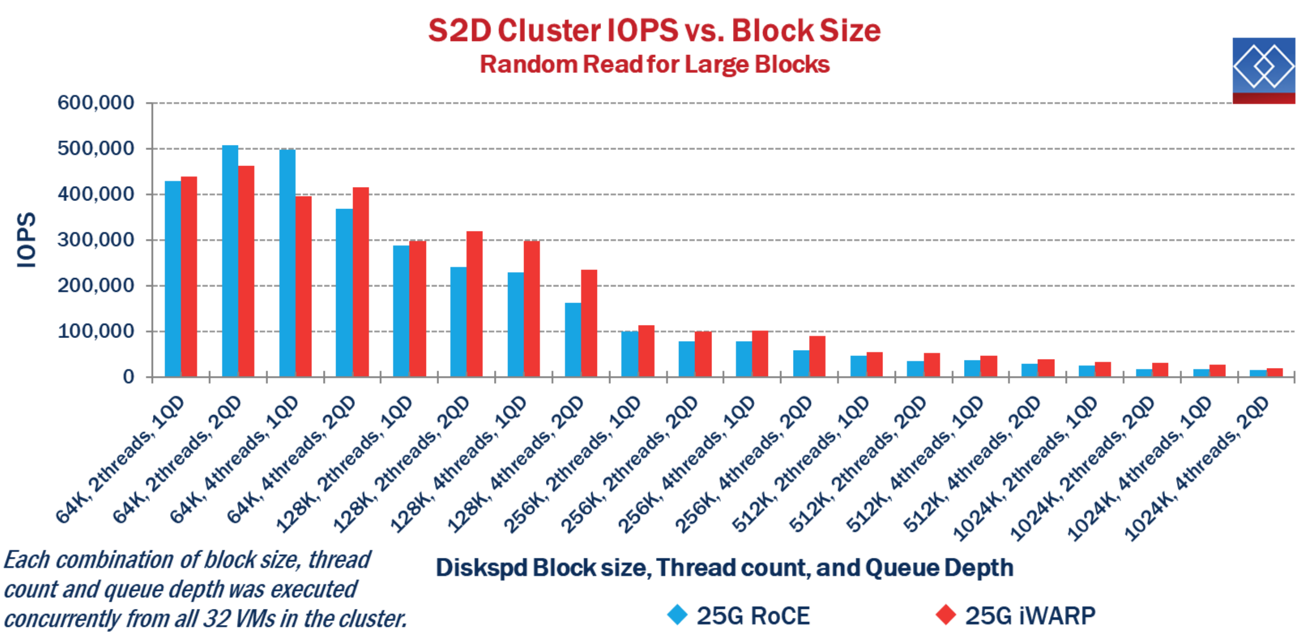

Each combination of block size, thread count and queue depth was executed concurrently from all 32 VMs in the cluster.

RoCE yielded higher total cluster throughput than iWARP for the smaller block workloads, but at larger block sizes iWARP delivered higher total cluster throughput.

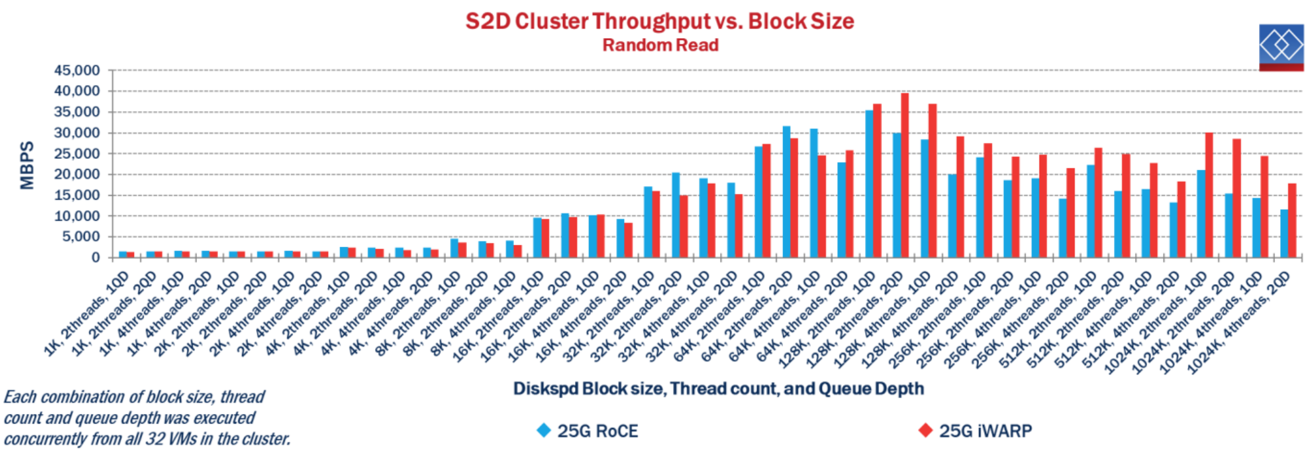

We can see that for both RoCE and iWARP, RDMA pays off, resulting in low processor utilization.

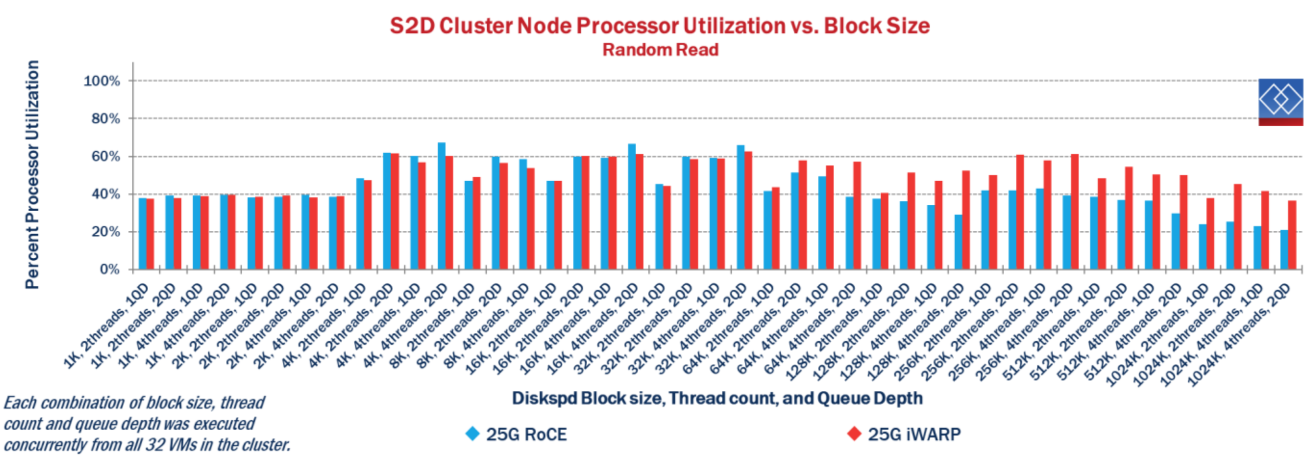

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Write Performance

Write performance will have higher latency than read due to the underlying storage media and traffic generated by the writes in a 3-way mirror.

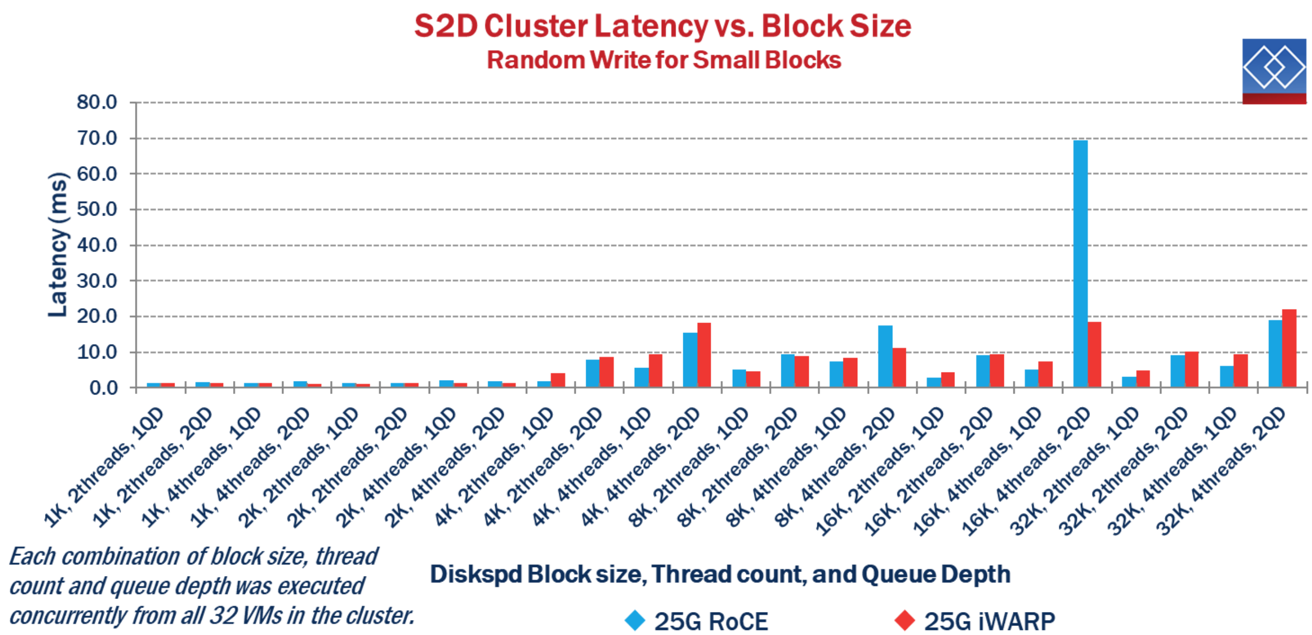

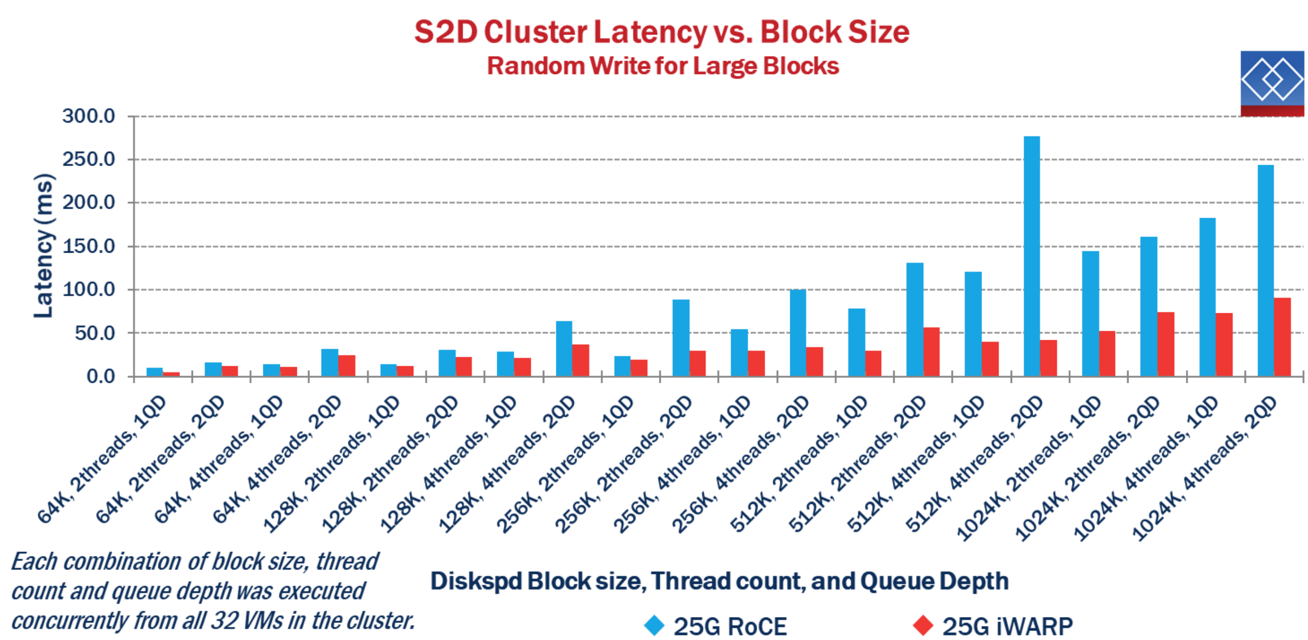

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

RoCE for large block sizes has higher latency than iWARP, but we do not see a corresponding increase in throughput. For very large block writes, iWARP provides more cluster throughput.

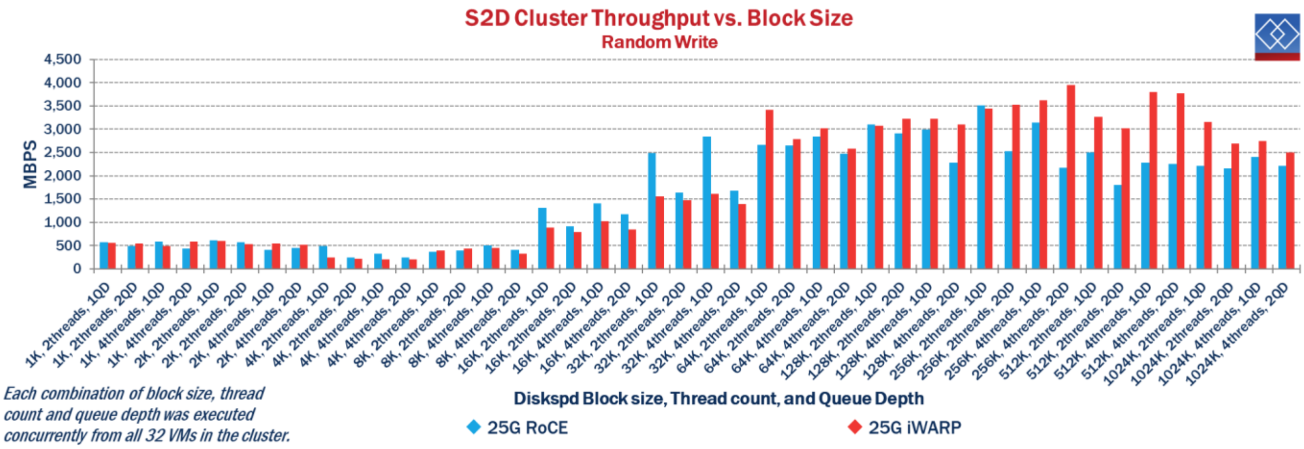

For both types of RDMA, total processor utilization is low. We see more processor utilization for iWARP for large block writes, most likely due to supporting the higher throughput that iWARP has in this case.

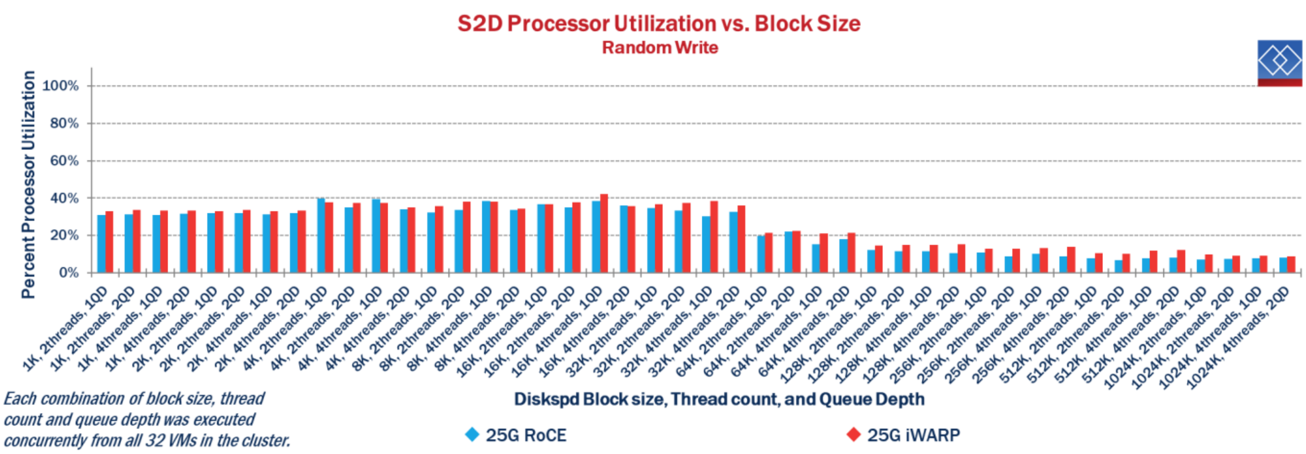

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

50% Read Performance

When performing 50% read, 50% write, the performance is a blend of what we have observed from our 100% read and 100% write testing. Latencies are higher for large block RoCE.

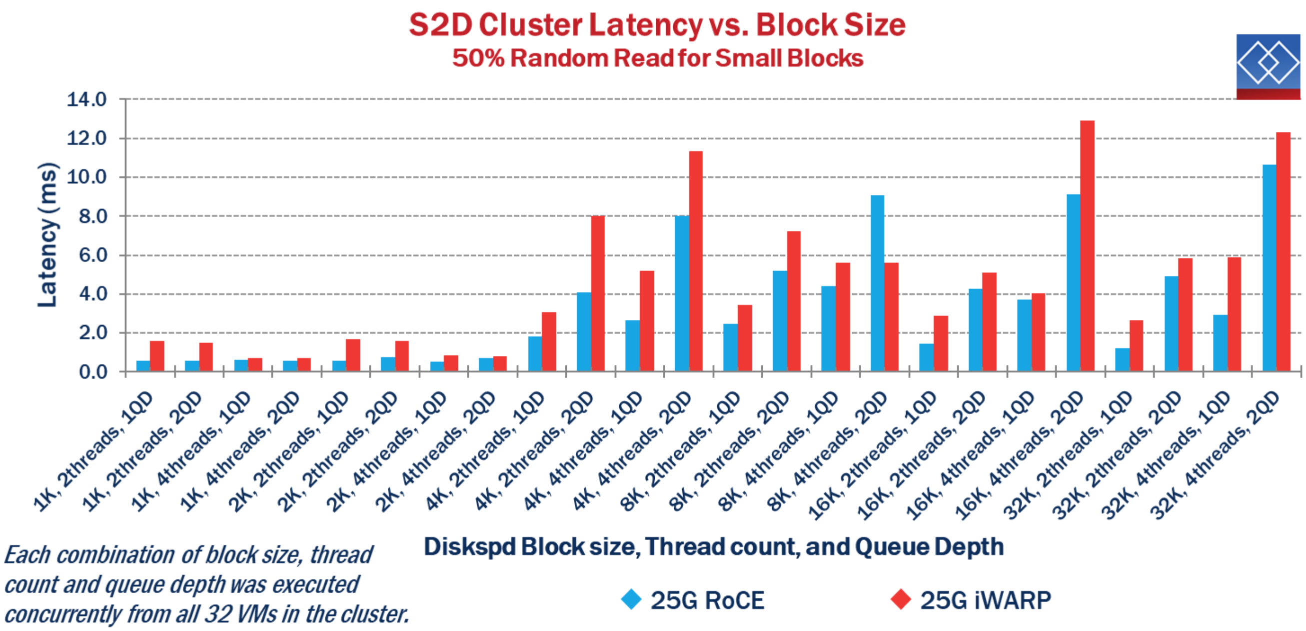

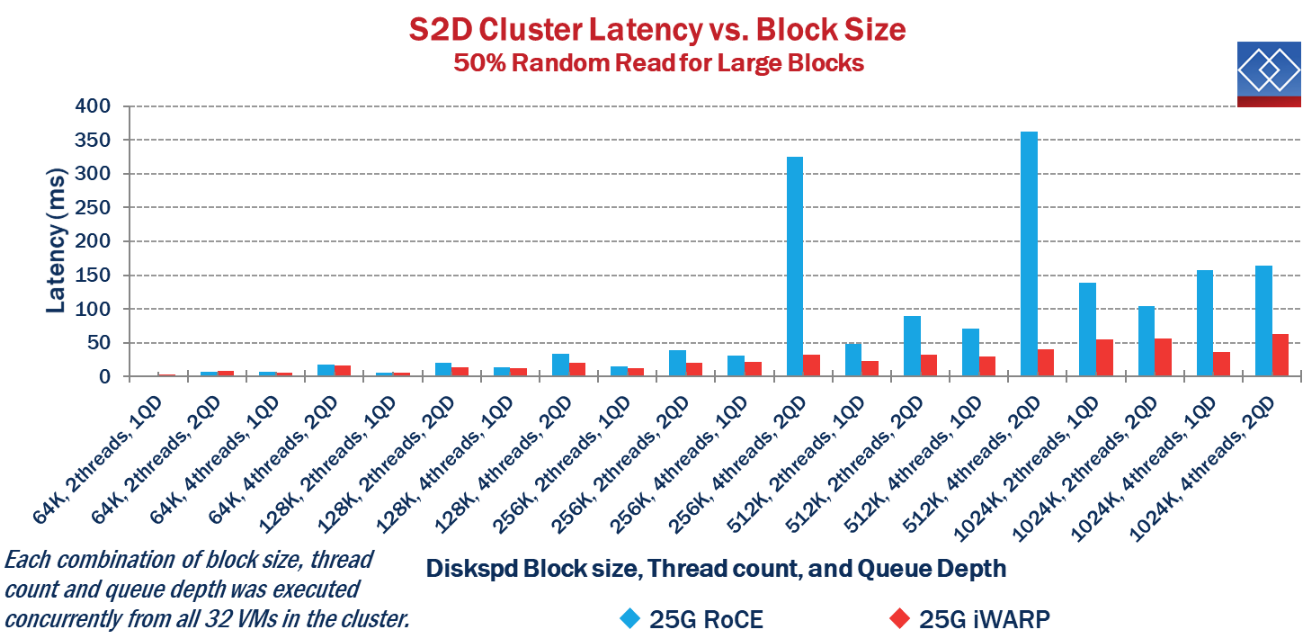

Similar to what we found in our 100% write testing, at large block sizes iWARP can handle more cluster throughput and uses a larger amount of processor to support that throughput.

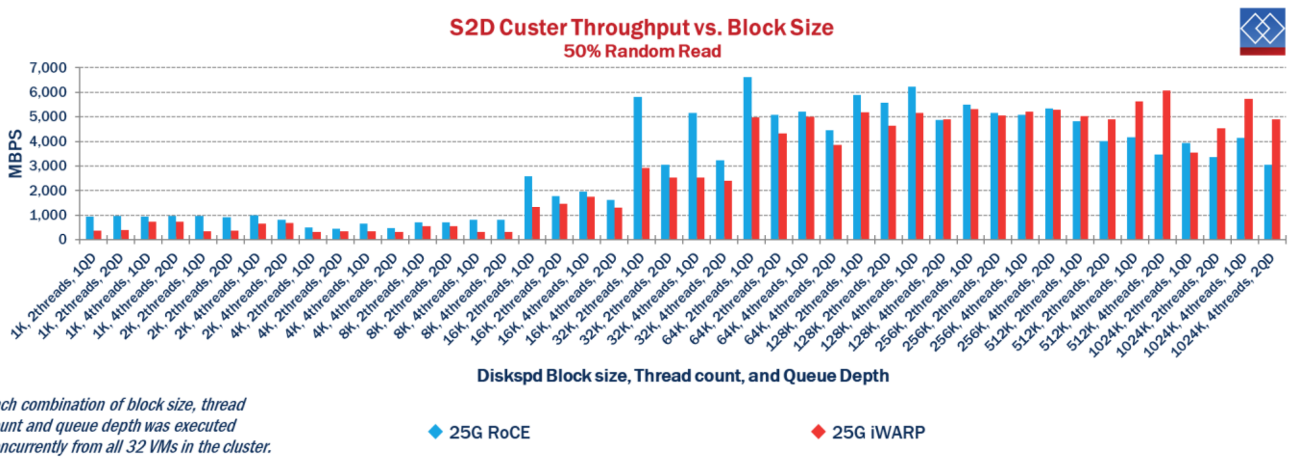

Again, processor utilization is very low due to our use of RDMA.

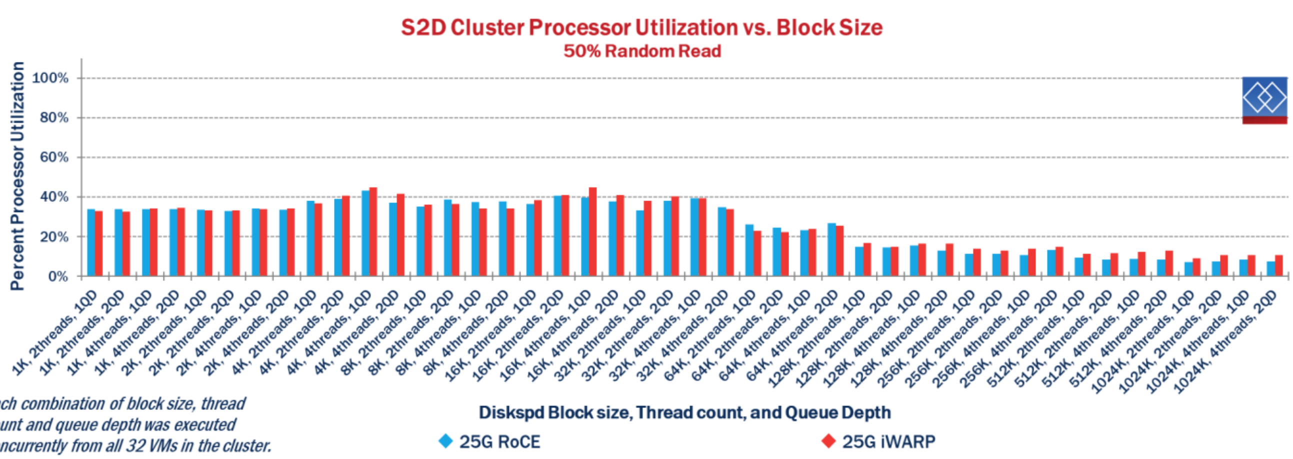

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Marvell FastLinQ 41000 Series 25GbE S2D Performance versus 10GbE S2D Performance – RoCE

The performance benefits of 25GbE translate to increased bandwidth for the S2D cluster.

Read Performance (RoCE)

25GbE had higher total cluster throughput than 10GbE for larger block workloads. As 25GbE begins to sustain higher throughput, the total processor utilization is higher.

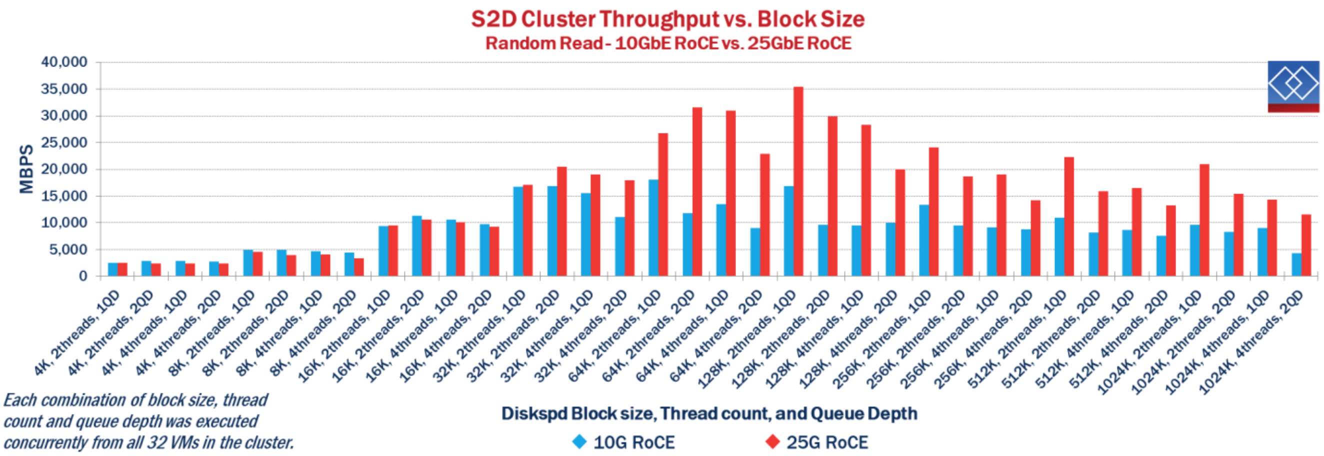

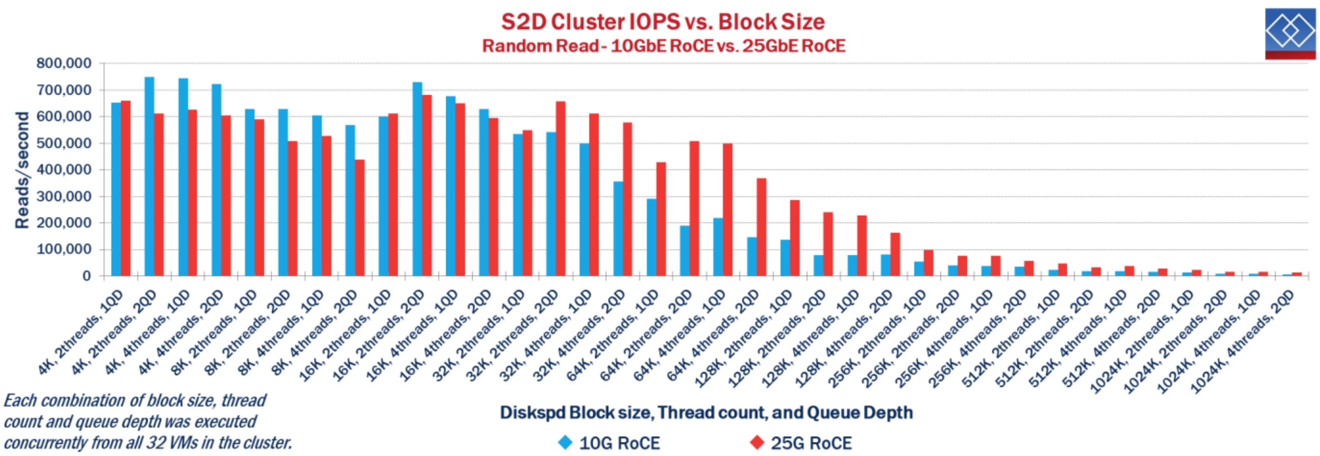

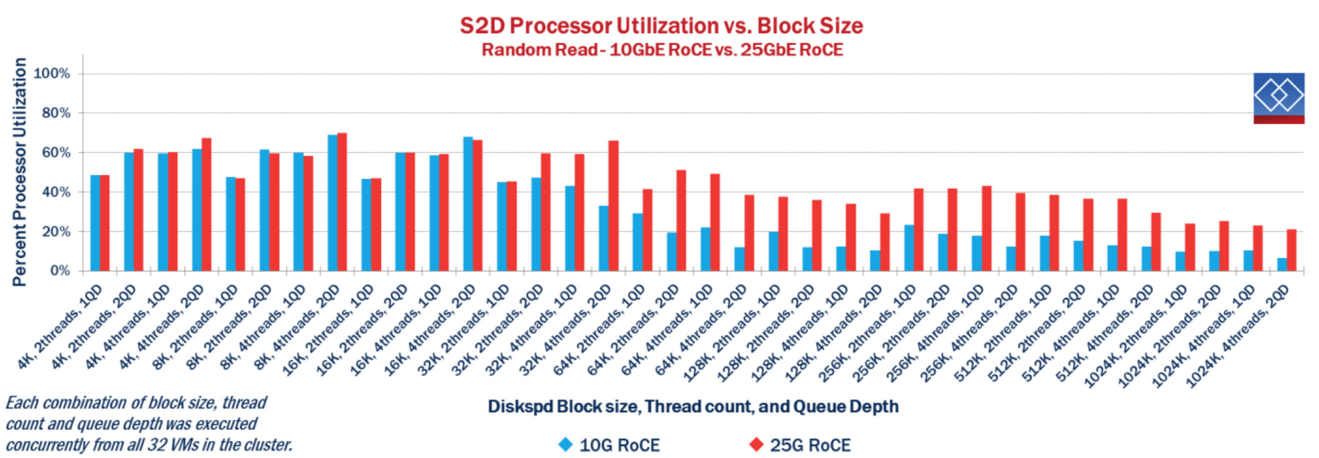

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Write Performance (RoCE)

For large block sizes, 25GbE RoCE once again sustained more throughput for writes.

25GbE processor utilization is slightly higher than 10GbE processor utilization for larger block writes.

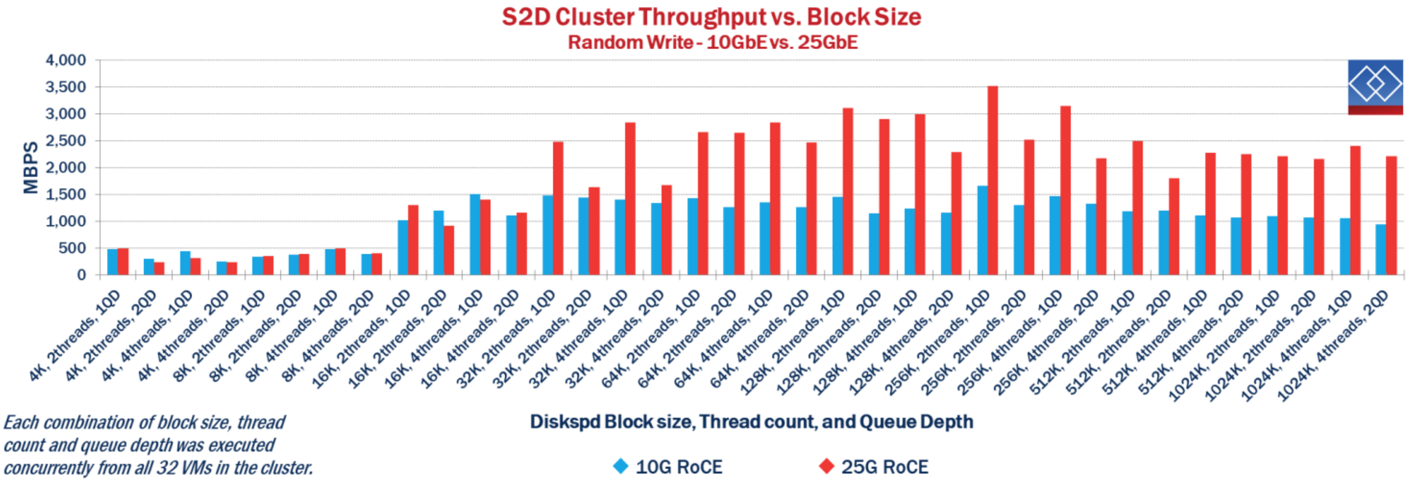

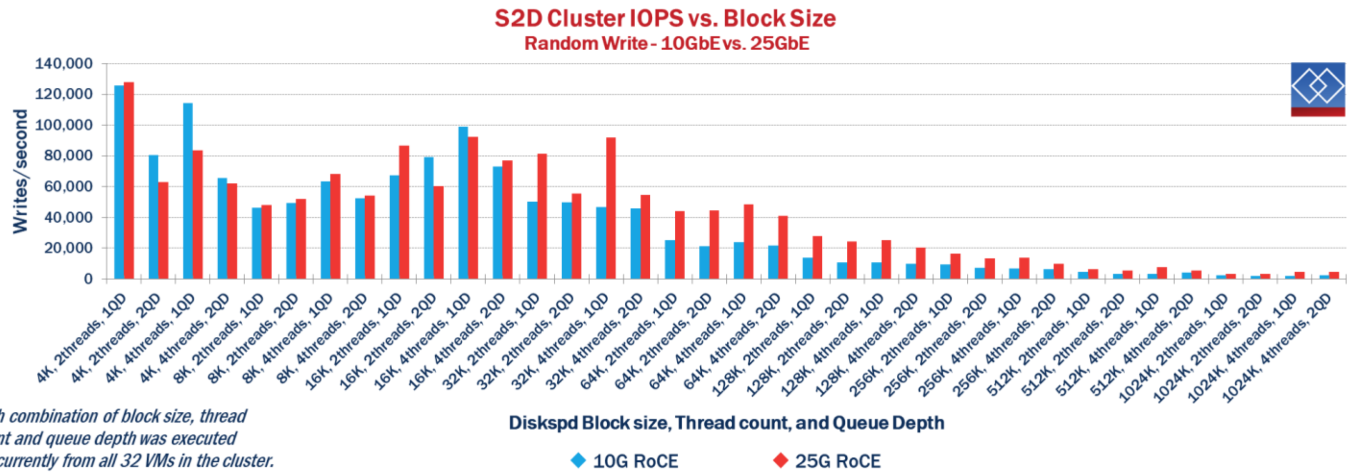

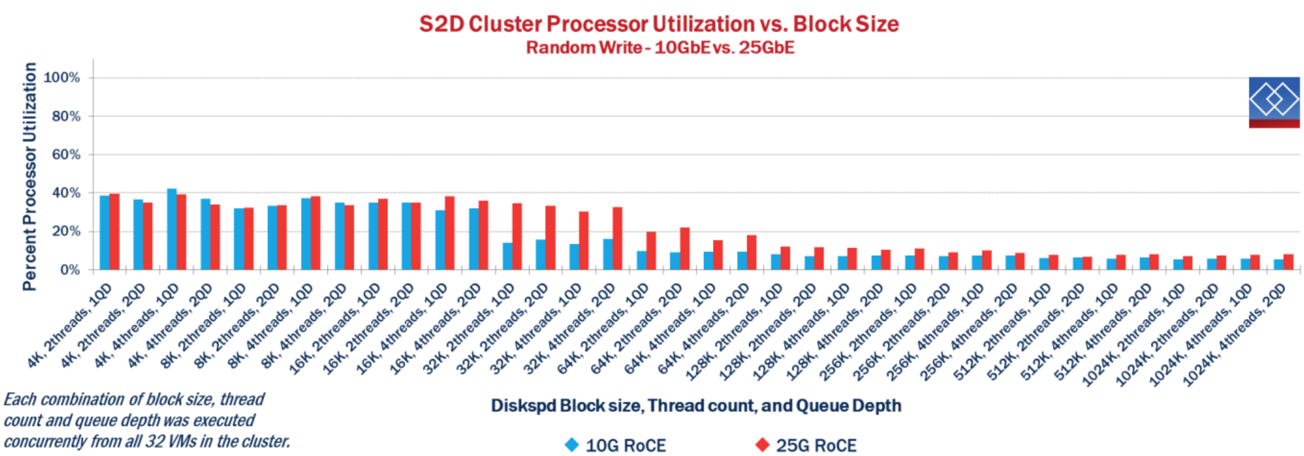

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

50% Read Performance (RoCE)

Once again, 25GbE can handle more throughput at larger block sizes than 10GbE can.

Especially at block sizes of 32KB and greater, many of the test cases yielded double or better throughput.

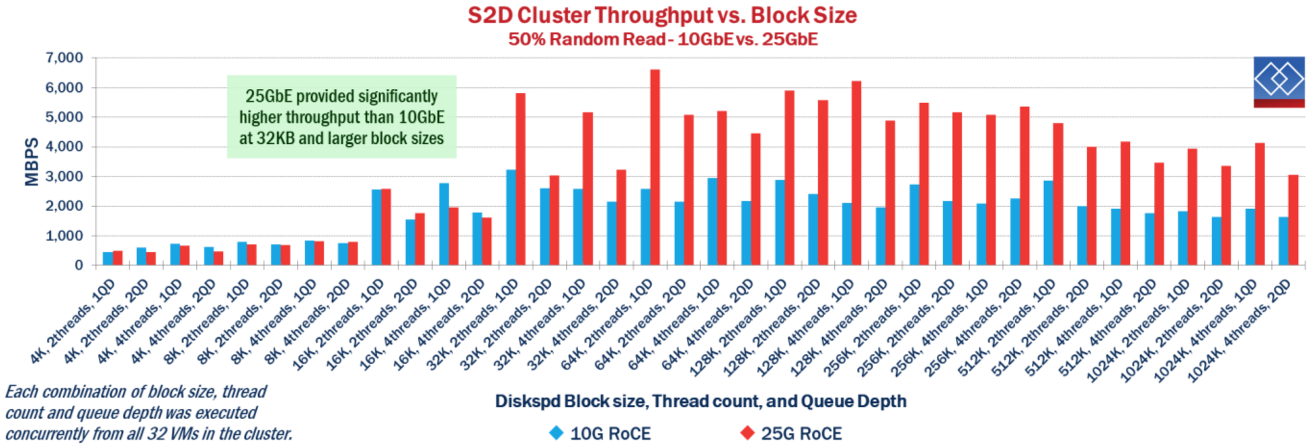

Processor utilization for 25GbE is slightly higher for large block sizes, a pattern that looks very similar to what was observed in the write test results.

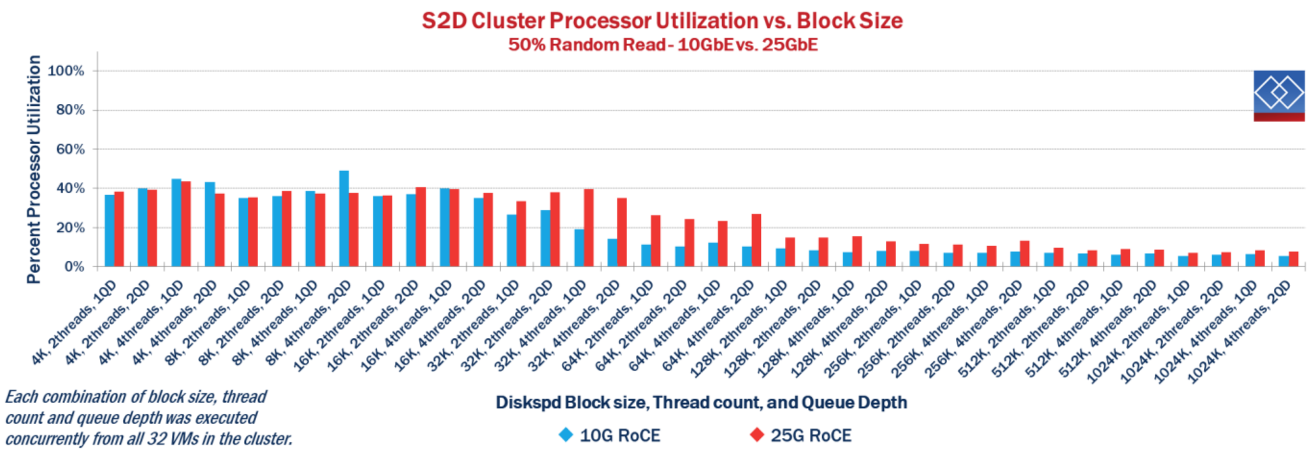

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Marvell FastLinQ 41000 Series 25GbE S2D Performance versus 10GbE S2D Performance – iWARP

The performance benefits of 25GbE translate to increased bandwidth for the S2D cluster.

Read Performance (iWARP)

25GbE had higher total cluster throughput than 10GbE for larger block workloads. As 25GbE begins to sustain higher throughput, the total processor utilization is higher.

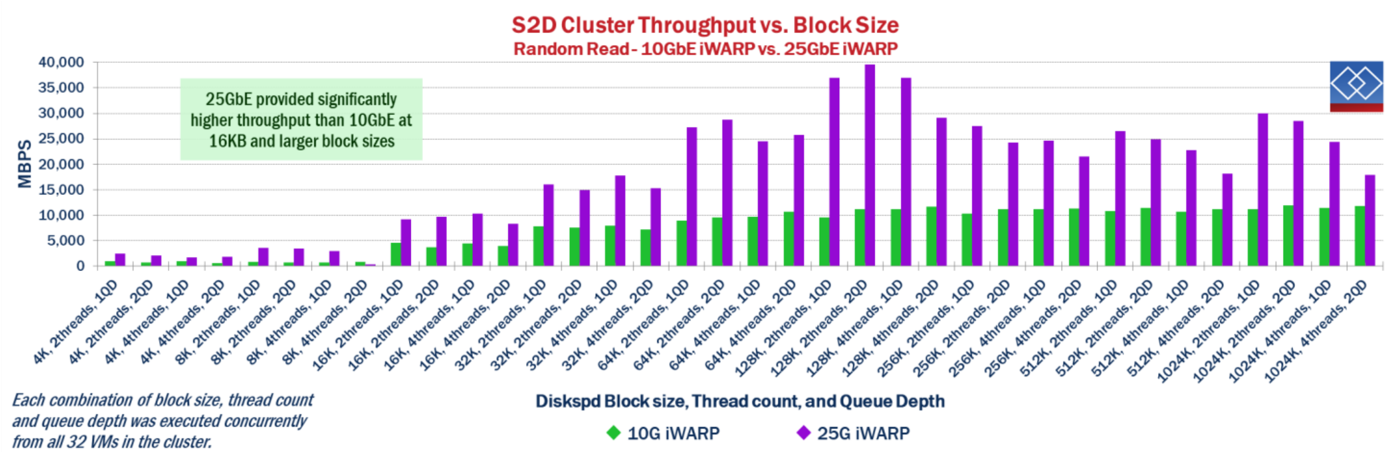

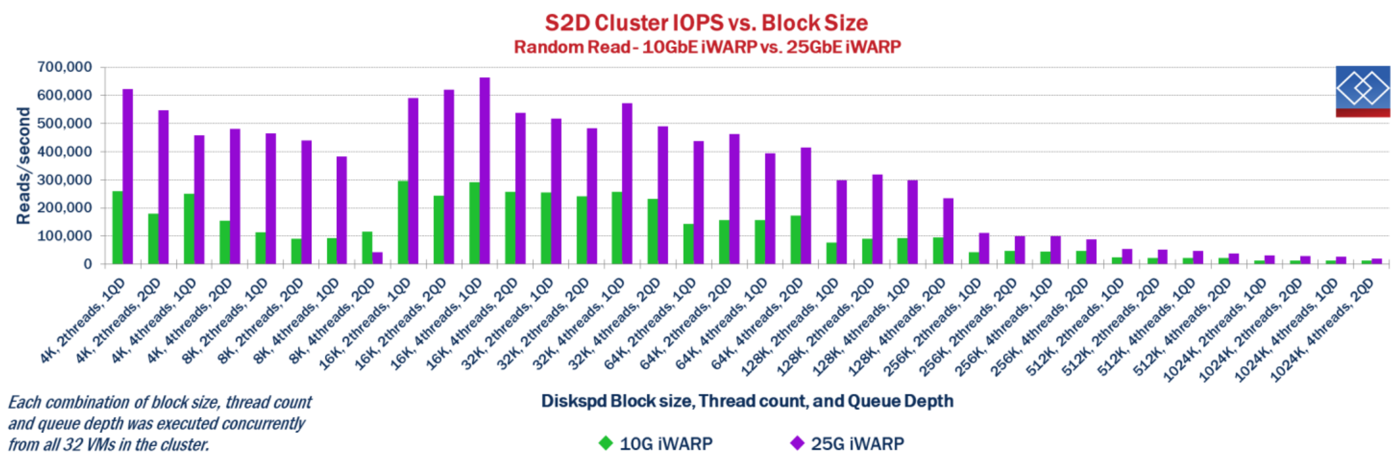

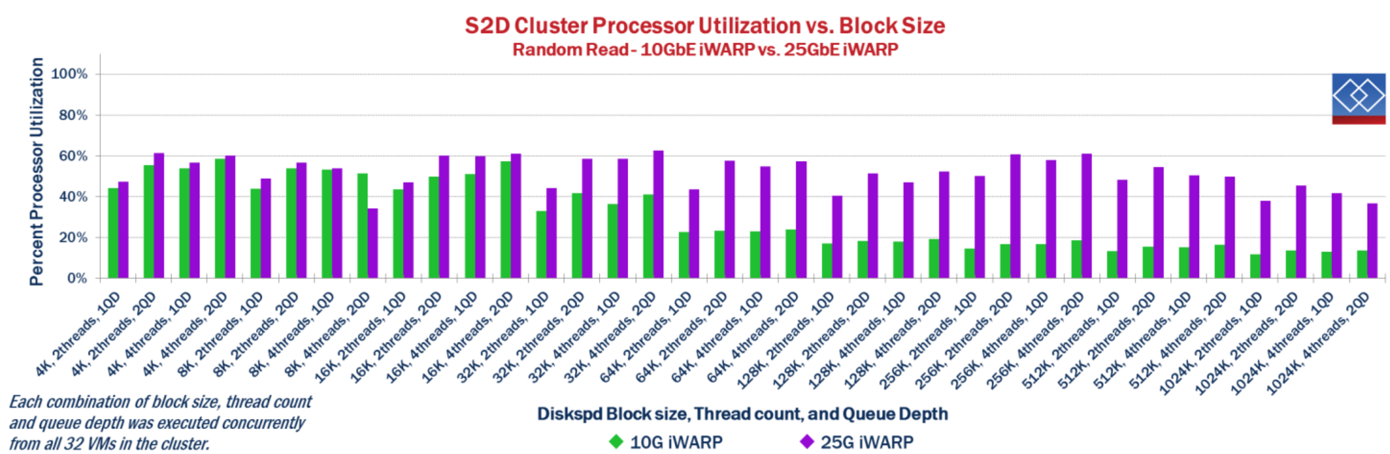

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Write Performance (iWARP)

For large block sizes, 25GbE iWARP once again sustained more throughput for writes.

Processor utilization for iWARP writes was nearly the same for 10GbE and 25GbE.

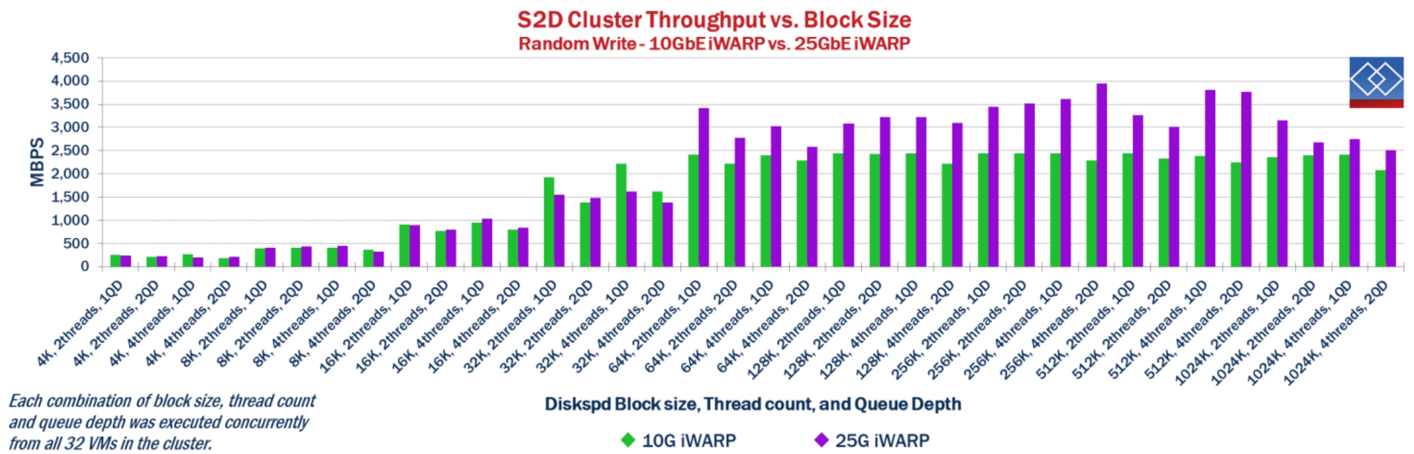

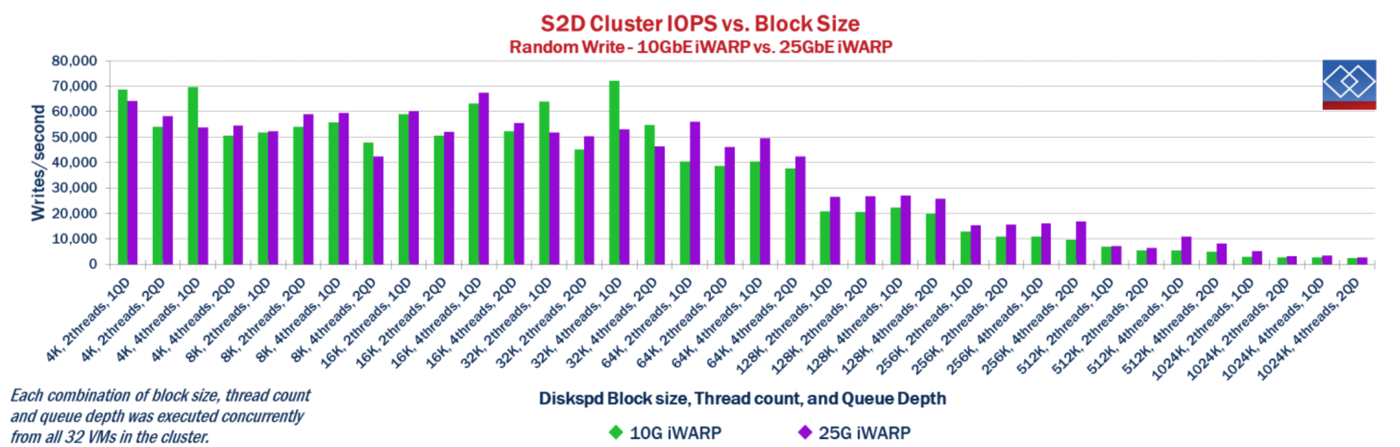

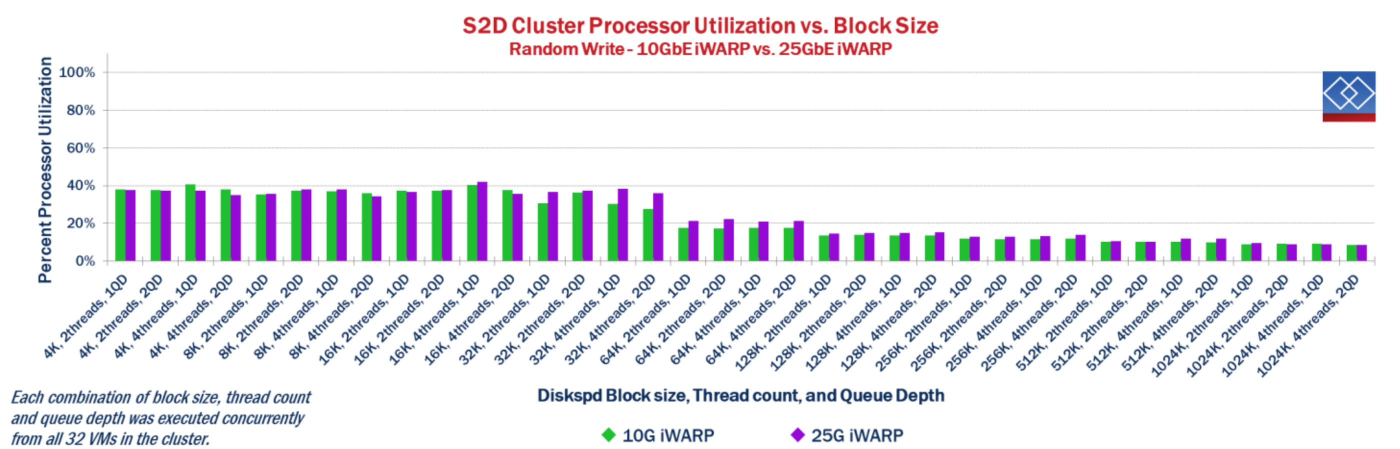

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

50% Read Performance (iWARP)

Once again, 25GbE can handle more throughput at larger block sizes than 10GbE can.

Throughput was consistently higher for 64GB sizes and higher for 25GbE iWARP. Processor utilization was very similar for iWARP for 10GbE and 25GbE.

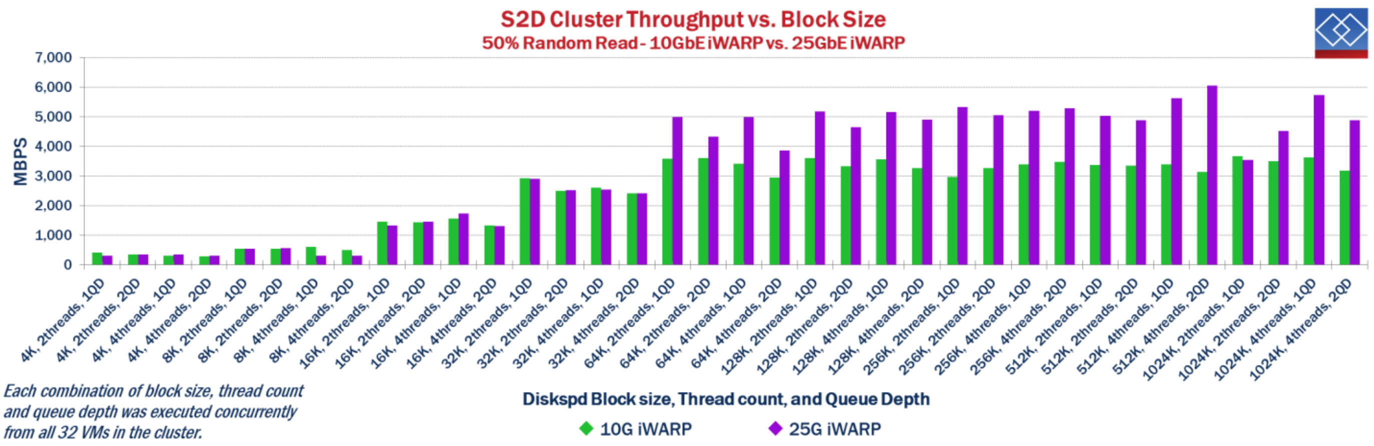

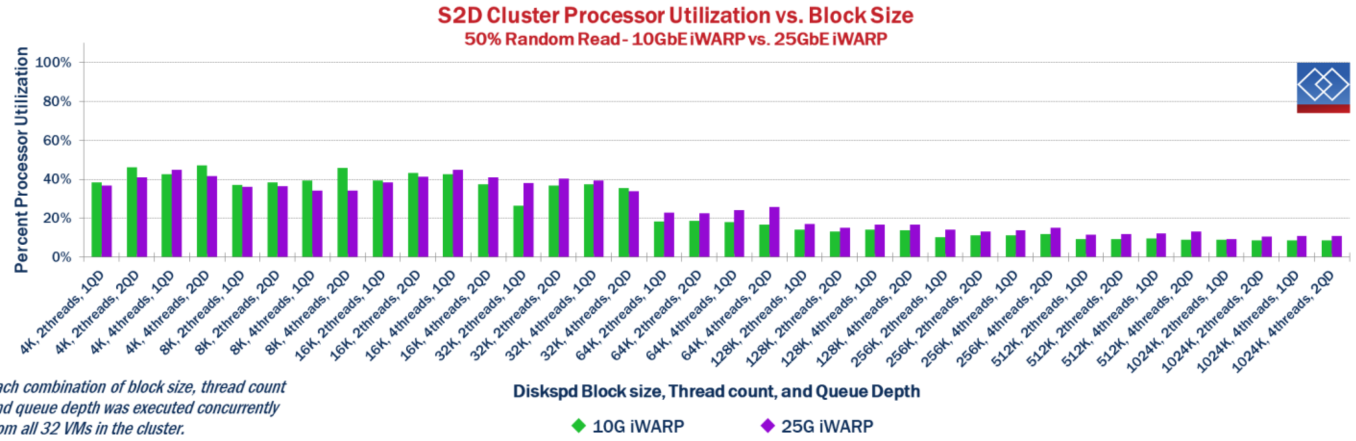

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

July 2019

Test Environment

Marvell FastLinQ 41000 Series Layer 2 Performance Tests

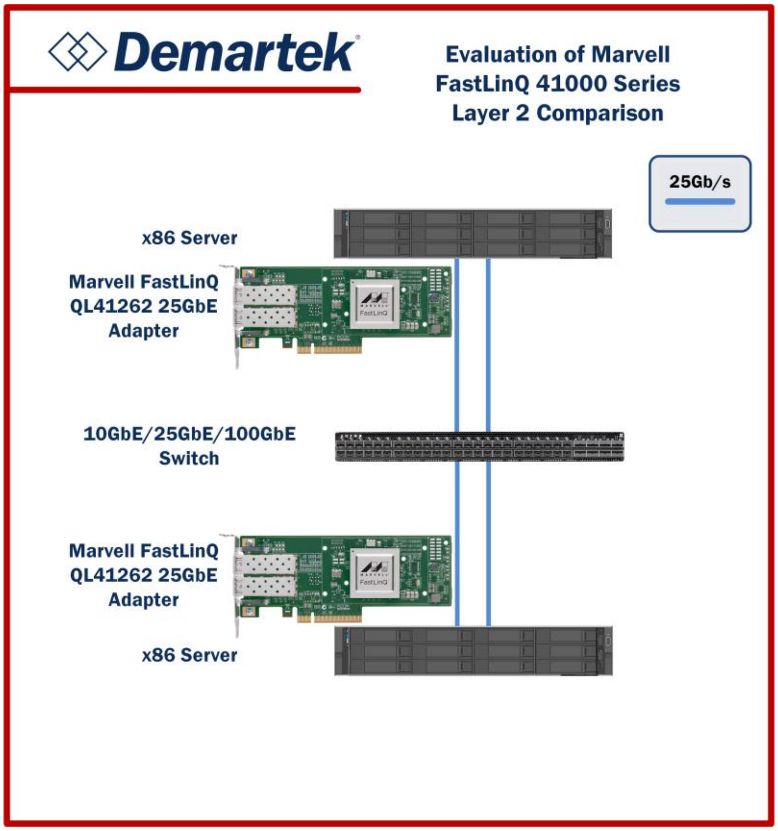

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

July 2019

Marvell FastLinQ 41000 Series iSCSI Offload vs Mellanox ConnectX-4 Lx Software iSCSI Tests

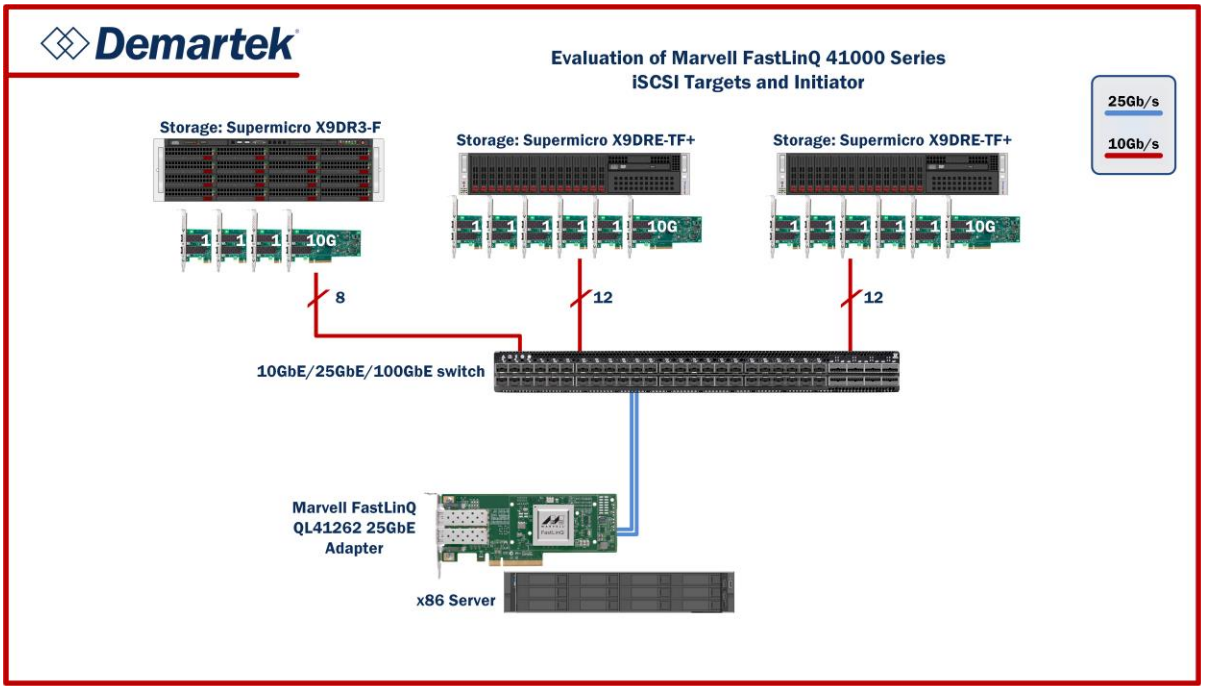

Marvell FastLinQ 41000 Series Storage Spaces Direct Tests

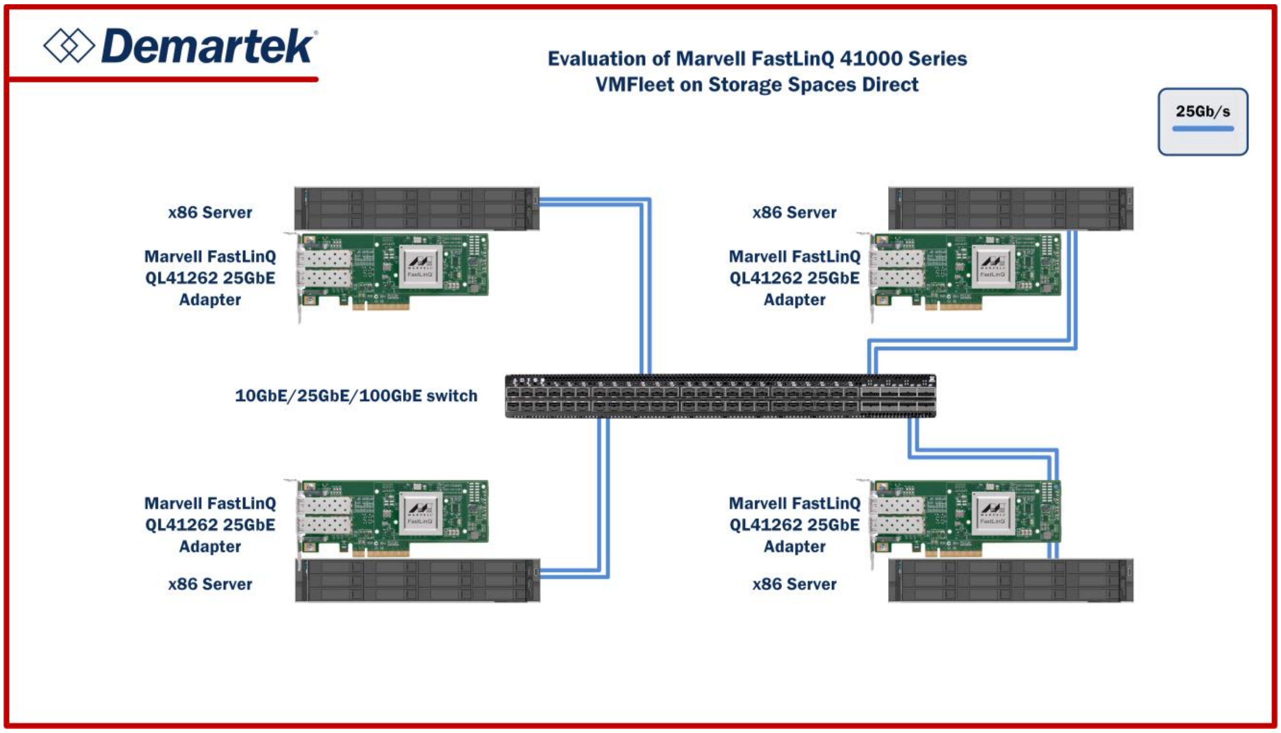

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Servers

> 2x Intel Xeon Gold 6130, 2.1GHz, 32 total cores, 64 total threads

> 192 GB Memory

> Microsoft Windows Server 2019 Datacenter Build 17723 (Windows Insider Preview)

Adapters

> Marvell FastLinQ 41000 Series, Boot Code 8.30.10.1, MBI 8.30.13, driver qede 8.30.12.0

OR

> Mellanox ConnectX-4 Lx when testing competitor software iSCSI, firmware version 14.20.1010 (MT_2420110034), driver mlx4_core 4.1-1.0.2

Switch

> Generic 10/25GbE/100GbE switch

iSCSI Storage Targets

2x Supermicro X9DRE-TF+

> 2x Intel Xeon E5-2690 v2, 3.0GHz, 20 total cores

> 256 GB Memory

> 6x Dual Port 10GbE NICs from multiple manufacturers

> RedHat Enterprise Linux 7.3

> Targetcli 2.1.fb41

1x Supermicro X9DR3-F

> 2x Intel Xeon E5-2690, 2.9GHz, 16 total cores

> 192 GB Memory

> 4x Dual Port 10GbE NICs from multiple manufacturers

> RedHat Enterprise Linux 7.3

> Targetcli 2.1.fb41

Marvell FastLinQ 41000 Series 25GbE Performance, iSCSI Offload Competitive Evaluation and Storage Spaces Direct Use Cases

Summary and Conclusion

The Marvell FastLinQ 41000 Series is an excellent choice for current generation servers with Intel Xeon Scalable processors, and 25GbE is the new standard for networking. While offering a broad choice of offloads and Universal RDMA, Marvell also achieves great performance.

> The Marvell FastLinQ 41000 Series achieved line rate bidirectional performance for buffer sizes of 2KB up to 1MB.

> The Marvell FastLinQ 41000 Series hardware iSCSI initiator achieved an average of 87% more IOPS than the Linux Software Initiator with Mellanox for unidirectional workloads. Results were similar for the bidirectional workloads.

> Overall, Marvell achieved line rate throughput on read and write while the Linux software initiator on Mellanox was only able to consistently achieve line rate on writes.

> The Marvell FastLinQ 41000 Series hardware iSCSI initiator consumed less processor than the Microsoft Windows software initiator on Mellanox, especially for larger block sizes.

> Marvell's ability to support both RoCE and iWARP RDMA provide customers with flexibility on their networking implementation which can reduce implementation time and/or simplify network complexity and management.

> Our S2D cluster built with Marvell FastLinQ 41000 Series Universal RDMA achieved on average 23,714 MBPS read throughput and 2,903 MBPS write throughput while using on average 41% of available cluster processor for reads and 13% of available cluster processor for writes.

> Our S2D cluster built with Marvell FastLinQ 41000 Series 25GbE Universal RDMA adapters achieved on average 126% more read throughput and 62% more write throughput when compared to the same cluster running 10GbE. Marvell's 25GbE Universal RDMA is necessary for S2D clusters to reach peak performance.