ESG Lab Review

The Performance Benefits of Fibre Channel Compared to iSCSI for All-flash Storage Arrays Supporting Enterprise Workloads

Date: July 2018 **Author:** Jack Poller, Senior Analyst

Abstract

This ESG Lab Review documents testing of Fibre Channel (FC) and iSCSI SAN fabrics. We focus on understanding the performance differences between the fabrics with all-flash storage arrays supporting enterprise workloads. Using three all-flash storage arrays, we measured IOPS, throughput, and latency, enabling direct comparison between each array's FC and iSCSI implementations. We also measured the impact of network congestion on the performance of each fabric.

The Challenges

The ever-increasing volume and velocity of data has made storage performance one of the top IT concerns. Indeed, according to ESG research, improved performance was, by far, the most often cited motivation for consideration or deployment of solid-state storage (see Figure 1).[1]

Figure 1. Factors Driving Solid-state Storage Usage/Consideration

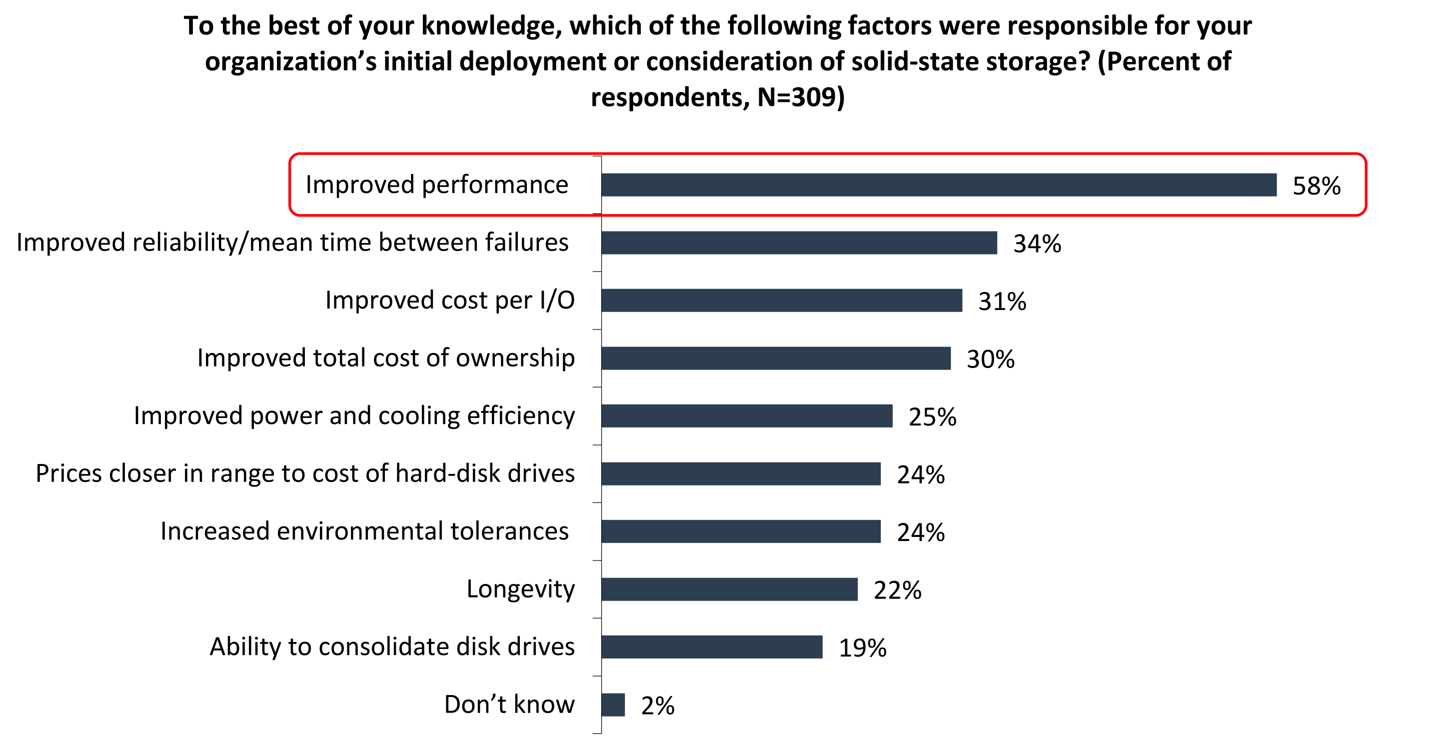

Source: Enterprise Strategy Group

Implementing high-performant storage infrastructures is often hampered by the increasing complexity of IT infrastructures. According to recent ESG research, more than two-thirds of surveyed organizations said that their IT environment has gotten

[1] Source: ESG Master Survey Results, *2017 General Storage Trends*, November 2017.

more complex in the last two years.[2] This complexity makes it more difficult to isolate and provision storage networks to minimize congestion and maximize performance. Thus, it's no surprise that organizations are seeking to identify the best storage network fabric for their all-flash storage arrays supporting enterprise workloads.

Storage Area Networks: Fibre Channel and iSCSI

High-performance all-flash storage arrays typically support both Fibre Channel and iSCSI as the network fabric between the servers and the storage arrays, and storage architects have the choice to select which connectivity model to use.

Fibre Channel

Purpose-built as network fabric for storage and standardized in 1994, Fibre Channel (FC) is a complete networking solution, defining both the physical network infrastructure and the data transport protocols. Features include:

- Lossless, congestion free systems—A credit-based flow control system ensures delivery of data as fast as the destination buffer can receive, without dropping frames or losing data.
- Multiple upper-layer protocols—Fibre Channel is transparent and autonomous to the protocol mapped over it, including SCSI, TCP/IP, ESCON, and NVMe.
- Multiple topologies—Fibre Channel supports point-to-point (2 ports) and switched fabric (2^{24} ports) topologies.
- Multiple speeds—Products are available supporting 8GFC, 16GFC, and 32GFC today.
- Security—Communication can be protected with access controls (port binding, zoning, and LUN masking), authentication, and encryption.
- Resiliency—Fibre Channel supports end-to-end and device-to-device flow control, multi-pathing, routing, and other features that provide load balancing, the ability to scale, self-healing, and rolling upgrades.
- Routing—Administrators can configure zoning to enable devices in two separate fabrics to communicate without merging the fabrics.

iSCSI

Standardized in 2004, iSCSI is a client-server SCSI transport layer protocol that defines how SCSI packets are transported over a TCP/IP network. Features include:

- Multiple physical networks—iSCSI can run over any existing TCP/IP infrastructure, including Fibre Channel, Ethernet, InfiniBand, and more. In practice, iSCSI typically runs over Ethernet networks, including production 1GbE, 10GbE, 25GbE, and 40GbE. Organizations can leverage existing network switches, physical plants, and personnel.
- Inherits TCP/IP features and functionality—iSCSI inherits all the benefits and properties of TCP/IP networks, including:
    - Security—Firewalls, VLANs, and other TCP/IP and network security solutions can protect iSCSI communication, while IPSEC provides in-flight encryption.
    - Routing—Clients and servers can communicate globally, without distance or network size limitations. iSCSI can be routed over the global Internet.
    - Resiliency—Multi-layer flow-control, multi-pathing, and discovery provide load balancing, the ability to scale, self-healing, and rolling upgrades.
- Support for virtual systems—iSCSI protocol can run in software, providing virtual systems with SAN support; there is no need to hand off data to an HBA, nor the need to emulate an iSCSI HBA.

[2] Source: ESG Master Survey Results, *2018 IT Spending Intentions Survey*, December 2017.

ESG Lab Tested

ESG began with an environment designed to benchmark the performance of Fibre Channel and iSCSI as network fabrics for all-flash storage arrays, as shown in Figure 2. The environment consisted of a Microsoft Windows Server 2016 with a 16GFC HBA and a 10G Ethernet NIC. The server was connected to both an Ethernet and a Fibre Channel switch. All three all-flash storage arrays were also connected to both the Ethernet and Fibre Channel switches.

As the environment was designed to observe the difference between Fibre Channel and iSCSI, three different leading all-flash storage arrays from three different vendors were deployed, reducing potential influence of driver, storage software, and storage hardware. Storage array A was a first-generation hybrid storage array, supporting both magnetic and solid-state storage. For this test, storage array A included only SSDs. Arrays B and C were second-generation, purpose-built all-flash storage arrays.

Figure 2. ESG Lab Test Bench

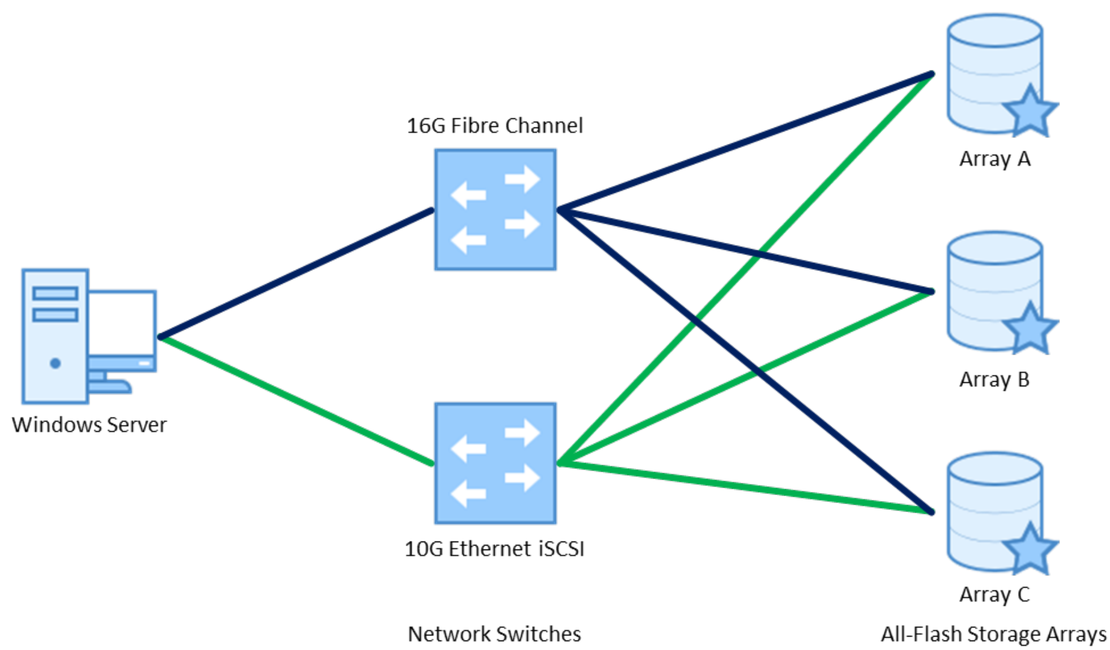

Source: Enterprise Strategy Group

ESG used Viavi Solutions' Medusa Labs Test Tools Suite (MLTT) to characterize the performance of the Fibre Channel and iSCSI networks. MLTT provides a comprehensive set of data integrity, benchmarking, and stress-test tools with precise and flexible control of storage network I/O and traffic.[3]

The testing was configured to perform only sequential reads and writes to obtain maximum performance from the storage arrays. ESG varied I/O block size from 8K to 128K, queue depths from 8 to 64 entries, and read/write ratio from 100% read to 100% write.

We focused on understanding the performance differences between Fibre Channel and iSCSI for typical enterprise use cases, such as OLTP. Thus, for this report:

- We report results for 70% read/30% write as representative of typical OLTP workloads.
- We report results for queue depths of 32 entries, which is the typical default.
- We report normalized results rather than absolute values.

The results we obtained for differing queue depths and read/write ratios are comparable to the results detailed in this report.

[3] https://www.viavisolutions.com/en-us/products/medusa-labs-test-tools-suite

Maximum Transactional Performance

To determine the maximum transactional performance of each fabric, ESG Lab ran a suite of MLTT tests, recording performance metrics of interest. Each test was designed to transfer as much data as possible as fast as possible to obtain the maximum performance of each storage array and fabric.

The peak number of I/O operations per second (IOPS) for each block size is shown in Figure 3. Each storage array's Fibre Channel performance is compared to its own iSCSI performance. Results are normalized to iSCSI—Fibre Channel IOPS are shown as a multiple of the same array's iSCSI IOPS.

Figure 3. All-flash Array IOPS

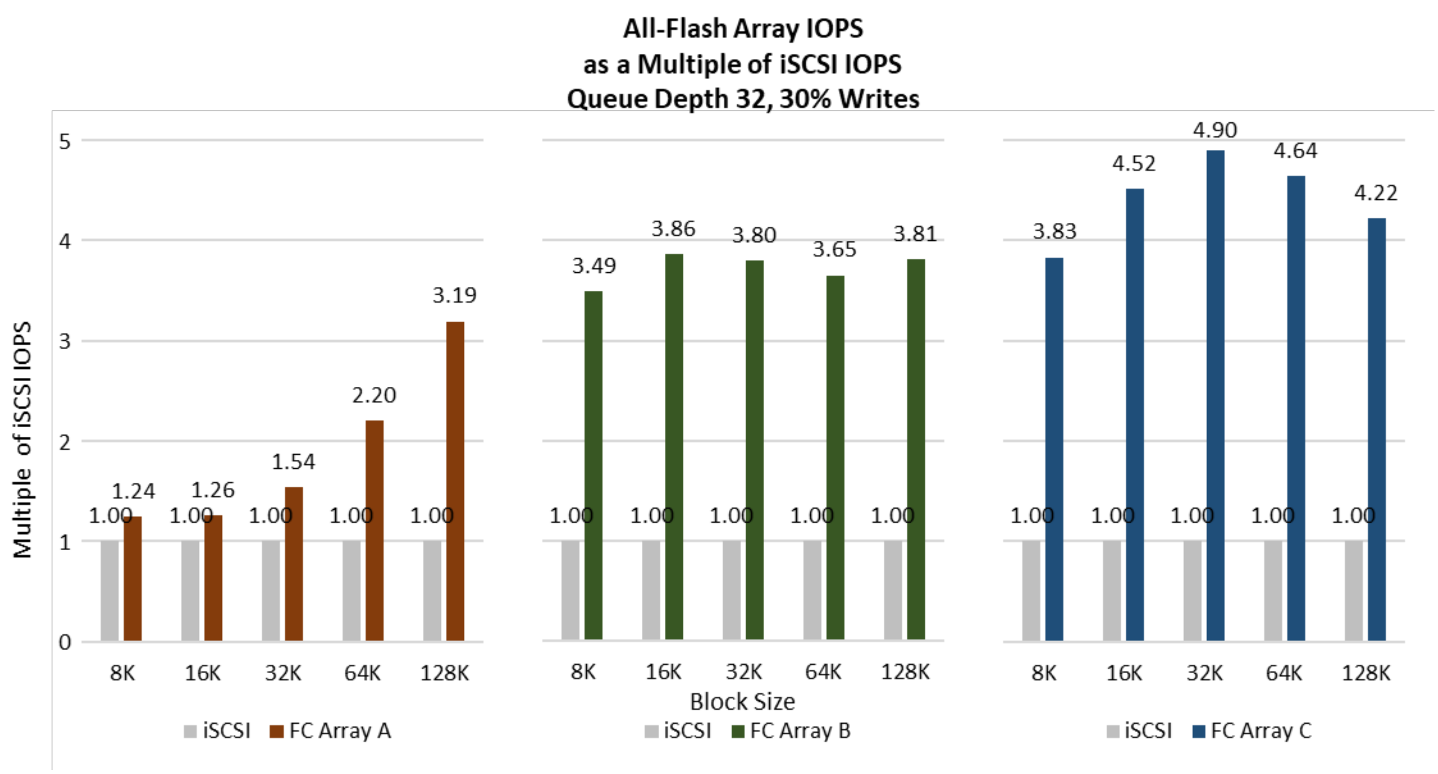

Source: Enterprise Strategy Group

What the Numbers Mean

- The first-generation all-flash storage array A sustained 25% more transactions with FC than iSCSI with small block sizes and more than triple the number of transactions with large block sizes.
- The second-generation all-flash storage arrays sustained between three and a half and almost five times more transactions than iSCSI depending on block size.
- While the Fibre Channel fabric provided 33% more throughput than the iSCSI fabric,[4] Fibre Channel was able to sustain almost five times as many transactions as iSCSI, demonstrating that Fibre Channel provides better transactional performance than iSCSI.

[4] The theoretical maximum throughput of 16GFC is 1,600 MB/second and 10GbE iSCSI is 1,200MB/second.

Maximum Throughput

Comparing IOPS between fabrics running at different speeds can provide misleading results. To provide a more reasonable comparison, we calculated the theoretical maximum amount of throughput of each fabric. The theoretical maximum throughput of 16GFC is 1,600 MB/second and 10GbE iSCSI is 1,200MB/second.

Figure 4 shows the results of the tests for maximum data throughput as a percentage of theoretical maximum throughput for the fabric for each block size. This representation shows how much of the pipe each fabric utilizes for each block size and enables direct comparisons between Fibre Channel and iSCSI for each array.

Figure 4. All-flash Array Throughput

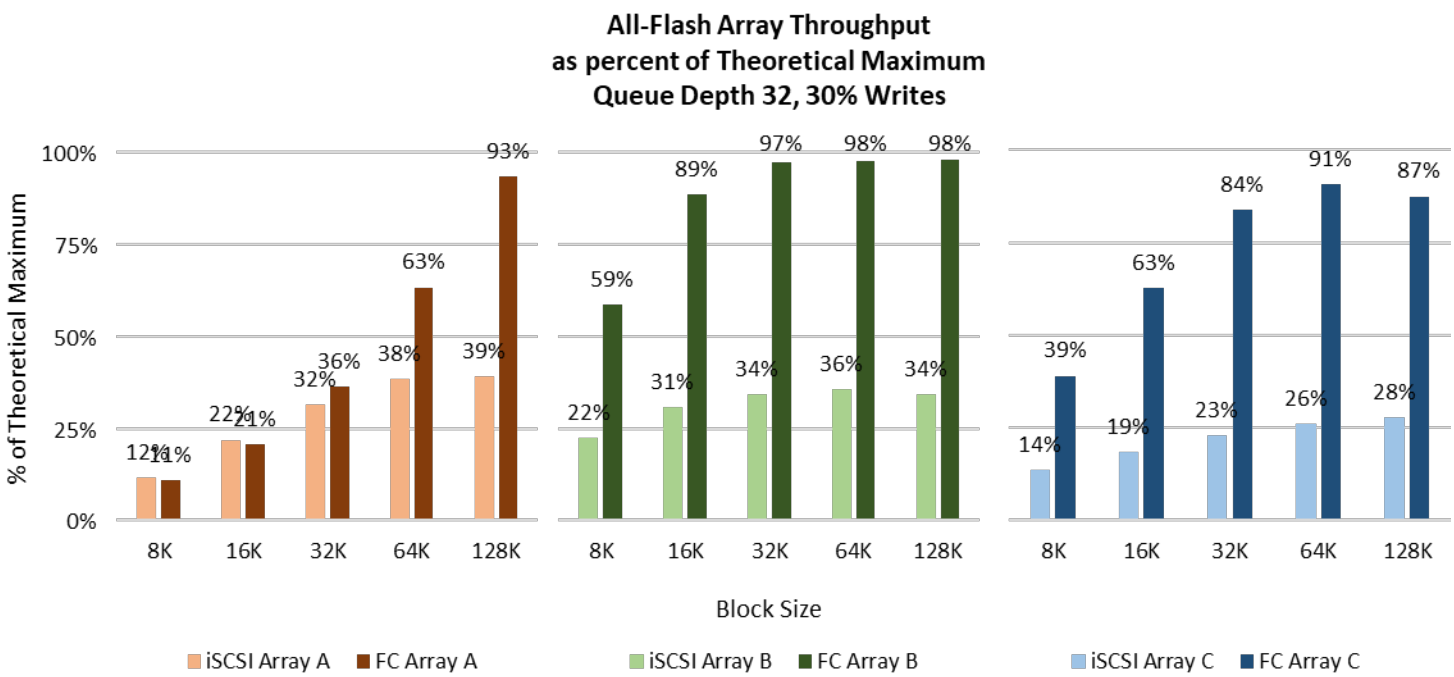

Source: Enterprise Strategy Group

What the Numbers Mean

- The first-generation all-flash storage array A provided nearly identical performance for both fabrics at small block sizes. With large block sizes, iSCSI consumed no more than 40% of throughput, while the Fibre Channel interface consumed almost all available throughput.
- The second-generation all-flash storage arrays consumed considerably more throughput using Fibre Channel than iSCSI. While the iSCSI fabrics consumed 14-40% of available throughput, the Fibre Channel fabrics consumed up to 98% of available throughput depending on block size. Because these arrays can drive much higher IOPS than Array A, they reach maximum throughput with smaller I/O sizes.
- At larger block sizes, second-generation all-flash storage array B consumed effectively the entire 16GFC pipe while only consuming at most 28% of the 10G iSCSI pipe.
- These results demonstrate that Fibre Channel is more efficient, utilizing more of the available throughput than iSCSI. Some implementations were able to consume the entire Fibre Channel throughput, yet no implementation consumed more than half of the iSCSI throughput.

Minimum Application Completion Time

MLTT calculates the average application completion time (also known as the round-trip time), a measure of latency, as the time between sending a transaction and receiving the acknowledgement on the Windows Server at the application layer.

The average application completion times for each block size are shown in Figure 5. Each storage array's Fibre Channel application completion time is compared to its own iSCSI application completion time. Results are normalized to Fibre Channel—iSCSI application completion time is shown as a multiple of Fibre Channel application completion time.

Figure 5. All-flash Array Application Latency

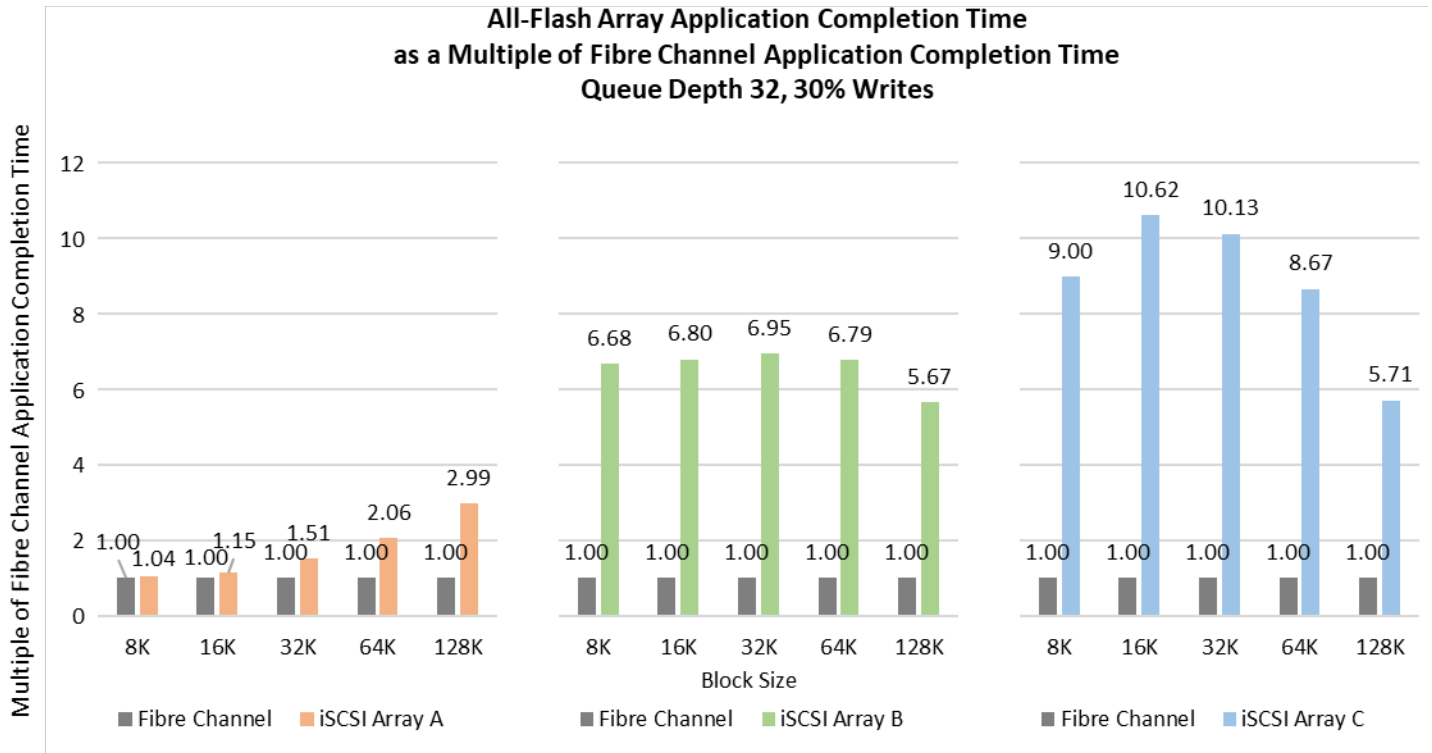

Source: Enterprise Strategy Group

What the Numbers Mean

- As with throughput, the first-generation all-flash storage array A (with FC latency up to 5 times slower than the second-generation arrays B and C) provided nearly identical results for both fabrics at small block sizes. With large block sizes, the round-trip time for iSCSI took up to 3 times longer than the Fibre Channel.
- Fibre Channel on the second-generation all-flash storage arrays B and C completed each storage transaction significantly faster than iSCSI, with Fibre Channel responding as much as 10 times faster.
- These results demonstrate that Fibre Channel is more efficient, completing transactions in much less time than iSCSI. Using iSCSI on an all-flash storage array may result in an application running up to ten times slower than when using Fibre Channel on the same array.

Congested Performance

To characterize how Fibre Channel and iSCSI perform with all-flash storage arrays supporting enterprise workloads, ESG modified the test bench to create congestion on each fabric, as shown in Figure 6. We added a Fibre Channel and Ethernet switch with redundant inter switch links (ISLs) between the two pairs of switches, simulating how systems are typically deployed in data center environments. We also added a traffic generator to the environment that could create Fibre Channel or Ethernet traffic, enabling us to simulate systems supporting enterprise workloads consuming fabric throughput.

Figure 6. ESG Lab Congestion Test Bench

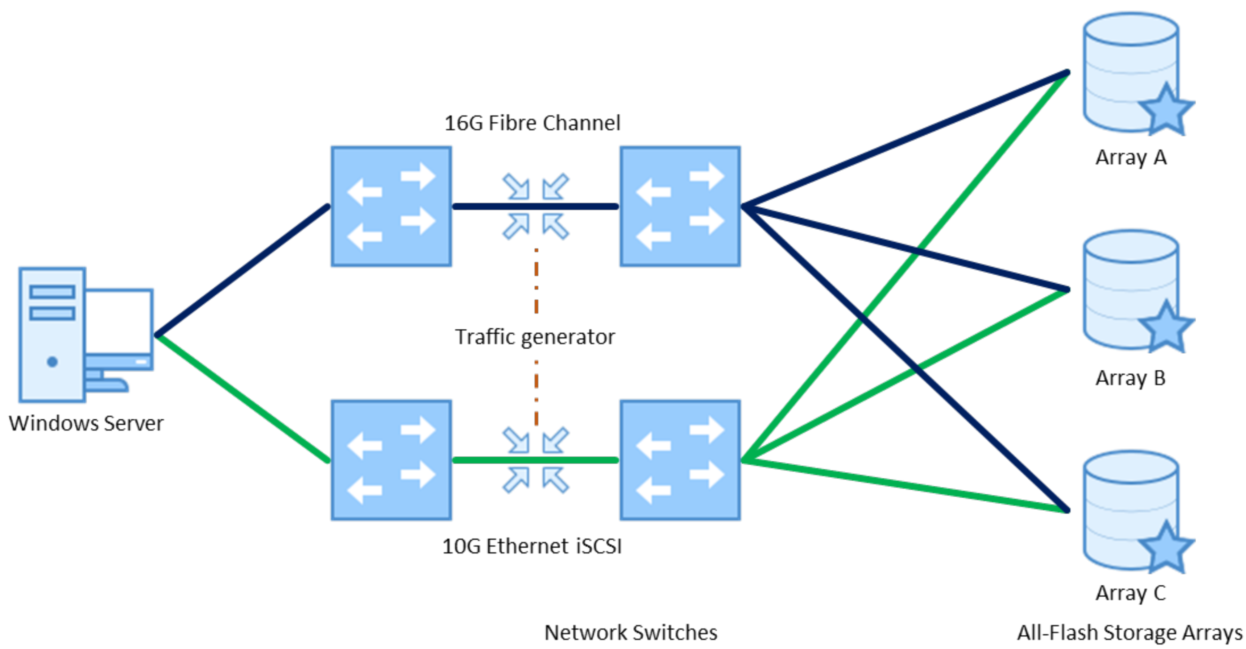

Source: Enterprise Strategy Group

We re-ran the tests with varying levels of background inter switch link utilization creating network contention, varying ISL utilization rates between 60% and 100%. Figure 7 shows the average throughput as a percentage of the maximum throughput achieved with no congestion in the fabric.

Figure 7. All-flash Array Throughput with Inter Switch Link Utilization

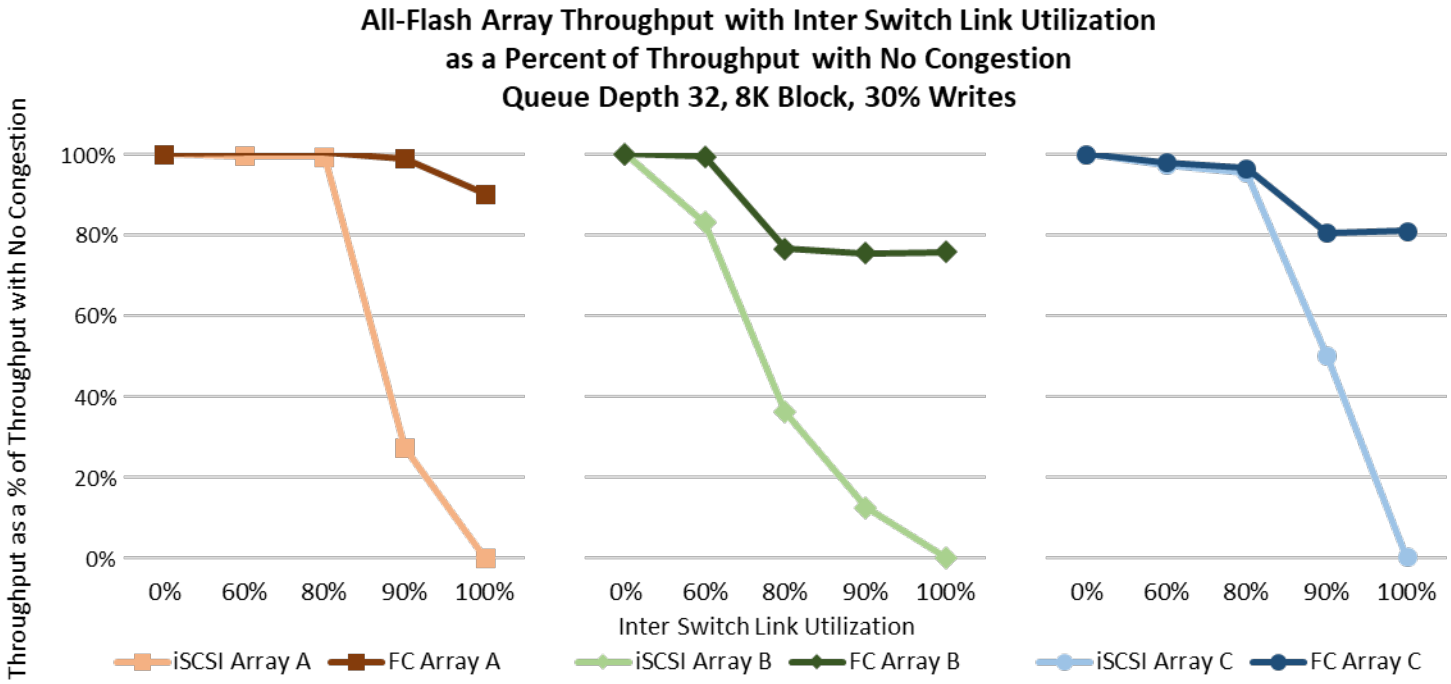

Source: Enterprise Strategy Group

What the Numbers Mean

- Because of the relatively low performance of the first-generation all-flash storage array A, no impact to throughput was observed below 80% ISL utilization. At 90% ISL utilization, iSCSI throughput dropped by two-thirds, and at 100% ISL utilization, iSCSI traffic effectively stopped. However, for this storage array, Fibre Channel was able to maintain the same throughput regardless of ISL utilization.
- The second-generation all-flash storage array B demonstrated consistently less iSCSI throughput as ISL utilization increased, dropping by two-thirds at 80% ISL utilization and by four-fifths at 90% ISL utilization. At 100% ISL utilization, iSCSI traffic effectively stopped. However, for this storage array, Fibre Channel throughput was reduced by 25% at 80% ISL utilization. The array was able to maintain the same Fibre Channel throughput at 80%, 90%, and 100% ISL utilization.
- Similar to the first-generation array A, the less-performant second-generation all-flash storage array C maintained the same iSCSI throughput as ISL utilization increased to 80%. However, throughput dropped to half at 90% ISL utilization, and at 100% ISL utilization, iSCSI traffic effectively stopped. For this storage array, like iSCSI, Fibre Channel maintained the same throughput as ISL utilization increased to 80%. Fibre Channel throughput dropped 20% between 80% and 90% ISL utilization, and this throughput was maintained at 100% ISL utilization.
- When network utilization increases, the performance of iSCSI degrades starting from 60% inter switch link utilization to practically a halt when ISL utilization is above 90%. This can make iSCSI unsuitable for applications requiring deterministic performance and response times while Fibre Channel performs consistently with high network utilization, in comparison.

Why does iSCSI throughput drop rapidly, coming to a halt, as the network becomes saturated, while Fibre Channel continues to provide significant throughput? Because the two fabrics have different flow control and congestion control mechanisms, and while TCP/IP, the underlying network protocol for iSCSI, allows for packets to be dropped, FC is a lossless network.

TCP/IP packet acknowledgements (acks) from the receiver to the sender include the *receive window*—the amount of buffer space available on the receiver, which tells the sender how much data can be in-flight between the two ends of the communication. However, the receive window only accounts for the receiver's buffer space, and not for any intermediary network nodes. Thus, as the network becomes congested, an intermediary node may run out of buffer space and start dropping packets, which requires retransmission.

Dropped packets and retransmissions can cause cascading congestion, as retransmissions consume more of the available throughput, leaving less throughput for new data blocks. In the worst case, as demonstrated by these tests, iSCSI transmission effectively stops, as the dropped packets and retransmissions consume all available throughput.

Fibre Channel operates on an end-to-end *buffer-to-buffer* accounting system. As the network starts up, each end of each link communicates the amount of buffer space available. A sender is responsible for tracking how much of the link's receiver buffer space the sender is consuming—each frame sent decrements the receiver buffer count, and each frame acknowledgement increments the receiver buffer count. A sender cannot send more data if the receiver buffer count is zero. Thus, as the network becomes congested, an intermediary node may run out of buffer space, causing the upstream sender to stop sending, which proceeds in turn all the way back to the originator of the communication.

FC's end-to-end flow control protocol includes intermediary nodes, which use fair share algorithms to ensure each sender gets their fair share of the available throughput as buffer space becomes available. Thus, as demonstrated by these tests, FC traffic continues to flow even as congestion approaches 100%.

Why This Matters

Storage is the foundation upon which organizations structure their modern data center architectures. More storage performance is required as infrastructures become more complex to support ever larger and more diverse workloads. The advent of the all-flash storage array has eliminated the bottleneck of the inefficient magnetic disk, improving access speed and throughput and exposing the storage network as the next performance bottleneck in the environment.

ESG Lab validated Fibre Channel outperformed iSCSI for Windows Servers using all-flash storage arrays. With no competing traffic, iSCSI consumed no more than 40% of maximum throughput while Fibre Channel was able to consume up to 98%. As the congestion increased, iSCSI throughput dropped, effectively stopping with a fully saturated network while Fibre Channel was able to maintain throughput at 75% (or higher) of a non-congested network with a fully saturated network.

These results indicate that Fibre Channel is more efficient and provides higher performance for networks with high levels of utilization.

The Bigger Truth

The ever-increasing volume and velocity of data, the shift to cloud architectures, and digital transformation initiatives are putting increasing demands on IT infrastructures. Storage environments must provide data protection, support for rapid data growth rates, automation, orchestration, and maximum performance. To meet these requirements while struggling to meet the omnipresent mandate of doing more with less, storage system architects frequently deploy all-flash storage arrays for primary storage.

The increased performance of all-flash storage arrays has put the spotlight on the storage network fabric as a factor influencing performance, especially in congested networks supporting enterprise workloads. Performance testing of iSCSI and Fibre Channel—the two predominant storage network fabrics—demonstrated the advantages of Fibre Channel.

In ESG Lab testing in a controlled environment, Fibre Channel consumed 98% of the maximum throughput in an isolated fabric, and delivered 75% of the maximum in a 100% saturated network. Conversely, iSCSI was only able to consume 40% of maximum throughput in an isolated environment. When the network was saturated, iSCSI traffic was effectively stopped as all throughput was consumed by dropped packets and retransmissions. Fibre Channel also demonstrated lower latency, with packet round-trip times taking as much as ten times longer for iSCSI than Fibre Channel.

Organizations seeking to maximize the performance of their all-flash storage systems supporting enterprise workloads should consider Fibre Channel for their storage network fabric for best ROI. Due to the many variables in each production data center environment, ESG recommends that you conduct your own testing to demonstrate Fibre Channel's performance benefits compared to iSCSI for your environment.

www.esg-global.com contact@esg-global.com P.508.482.0188