Evaluation Report: HP Blade Server and HP MSA 16GFC Storage Evaluation

Evaluation report prepared under contract with HP

Executive Summary

The computing industry is experiencing an increasing demand for storage performance and bandwidth due to increases in virtual machine density, increasing demands for application performance and continual data growth. Fibre Channel Storage Area Networks (SANs) carry the bulk of storage traffic in the enterprise data center and are beginning to feel the stresses of these increased demands.

As a result, enterprises are finding themselves constrained by the available bandwidth between the servers and storage, or foresee a constraint as they observe their growing data consumption patterns. HP provides an end-to-end Gen 5 (16Gb) Fibre Channel solution with the HP c7000 BladeSystem Enclosure, HP ProLiant BL460c Gen8 servers, 16GFC HBAs, embedded 16GFC switches and HP MSA 2040 16GFC storage system. This solution addresses these increasing demands on storage performance by providing double the bandwidth of previous generation Fibre Channel storage infrastructure.

Demartek deployed the complete end-to-end 16GFC solution described above, and ran a read-intensive data warehouse workload to determine whether this type of workload could take advantage of the increased bandwidth and performance that Gen 5 Fibre Channel provides. We repeated the database workload test with previous-generation HP 8Gb Fibre Channel HBA and SAN switch products then compared the results.

Key Findings

For the database workload, the 16Gb Fibre Channel end-to-end solution exceeded the performance of the 8Gb Fibre Channel end-to-end solution. The 16Gb Fibre Channel solution provided the additional bandwidth needed to complete the job in significantly less time.

HP end-to-end 16GFC results:

- The real database workload was completed 37% faster than the 8Gb infrastructure, reducing the time to complete the 6-user workload by approximately four hours.

- Reduced latency by approximately 59%.
- For highly virtualized environments and data intensive DB applications, 8Gb infrastructure can saturate and become the bottleneck for maximizing I/O performance.

The Need for More Bandwidth

Today's datacenters face a variety of challenges brought on by seemingly insatiable demands on server and storage infrastructure. At the same time, new technologies are being introduced that offer both challenges and possible solutions to meet these growing challenges.

The following is a summary of factors that Demartek has identified as driving the need for more storage networking bandwidth as provided by Gen 5 Fibre Channel infrastructure.

VM Density

When Demartek presents next-generation storage networking technologies at various end-user industry events, we ask the primarily technical audience a few questions about their environments. Among the responses are that VM density has been increasing over the last few years, with higher numbers of guest operating systems running on one physical server than in the past. We expect this trend to continue. VM density growth, along with increasing physical densities of newer blade servers such as the HP c-Class, allow for very large numbers of virtual machines in one blade chassis, or in one rack full of blade chassis.

8Gb Fibre Channel Saturation

During the past two years, when we asked the end-users in our audiences about saturation of Fibre Channel links, we consistently heard from users who indicated that they had saturated their 8Gb Fibre Channel infrastructure and needed something faster. The applications consistently identified as needing this higher bandwidth were database applications, regardless of the brand of database. These include single database instances running on physical hardware, multiple database instances running on physical hardware and multiple database instances running in virtual machines (VMs). These users are generally looking for something compatible with their existing infrastructure but that provides higher bandwidth to meet their growing demands.

SSD Technology

Solid State Disk (SSD) technology is another driver of bandwidth growth. Although relatively early in the deployment cycles, we have found that those who deploy any form of SSD technology in the enterprise have experienced significant storage performance improvements. Many of these SSD deployments are in SAN environments, which drive up storage networking bandwidth consumption. Based on comments from users and many of the tests we have performed in our own lab, we have concluded that faster storage networking technology such as Gen 5 Fibre Channel—including the HP MSA 2040 storage system—is well-suited to SSD technology.

New HP Servers

In the spring of 2012, HP began to introduce rack and blade servers that support the Intel® Xeon® E5-2600 processors (code name "Romley") and in the autumn of 2013 servers that include the Intel® Xeon® E5-2600 V2 processors (code name "Ivy Bridge"). These servers support PCI Express (PCIe) 3.0. These servers provide not only higher numbers of cores and performance improvements in processor power but also provide significant increases in I/O throughput.

PCIe 3.0 doubles the maximum possible I/O rates to 1 GBps (gigabyte per second) per lane from the previous generation. PCIe 3.0 also doubles the maximum number of lanes available up to 40 PCIe lanes per processor as compared to the previous generation. As a result, the total I/O bandwidth available in one of these new servers is approximately quadruple that of the previous generation of servers.

Bandwidth Growth Summary

When we discuss storage networks with enterprise users, we find that Fibre Channel is still the dominant storage interface in large-scale data centers, and is expected to remain dominant as a SAN interface for the foreseeable future.

When using storage-intensive applications like backup/restore, database transactions, virtualization and rich media, there is clearly a need for higher storage networking bandwidth and performance. The improved I/O performance of Gen 5 Fibre Channel enables faster storage and retrieval of data. For those enterprises that don't believe that they need this higher performance yet, now is the time to start planning for these eventualities.

Gen 5 (16Gb) Fibre Channel

Generation 5 Fibre Channel provides not only doubles the throughput from the previous generation, but has other benefits. Some of these benefits are directly related to the speed increase, such as a reduced number of links needed to achieve the same bandwidth, reduced power consumption needed to achieve the same bandwidth and fewer cables to manage. In addition, the higher speed allows fabrics to be connected with fewer inter-switch links (ISLs), which is especially helpful in large fabrics.

Gen 5 Fibre Channel includes 64b/66b encoding and transmitter training. These features improve link performance characteristics, provide optimal signal quality and more efficient communication at higher speeds than were previously used for Fibre Channel.

Table 1 – Fibre Channel Speed Characteristics

Speed	Throughput (Mbps)	Line Rate (Gbps)	Encoding	Transmitter training
1 GFC	100	1.0625	8b/10b	No
2 GFC	200	2.125	8b/10b	No
4 GFC	400	4.25	8b/10b	No
8 GFC	800	8.5	8b/10b	No
16 GFC	1600	14.025	64b/66b	Yes

Table 2 – Fiber Optic Cable Link Distance (in meters)

Speed	Multi-Mode				Single-Mode
	OM1	OM2	OM3	OM4	OS1
1 GFC	300	500	860	*	10,000
2 GFC	150	300	500	*	10,000
4 GFC	50	150	380	400	10,000
8 GFC	21	50	150	190	10,000
16 GFC	15	35	100	125	10,000

* The link distance for OM4 fiber optic cable has not been defined for these speeds.

HP BladeSystem c7000 and ProLiant BL460c Gen8 Servers

HP BladeSystem combines a custom-designed chassis with modular server, networking and storage components to create a fully functioning computing system, sometimes called a "mini-datacenter."

The BladeSystem c7000 Platinum Enclosure provides the power, cooling, and I/O infrastructure needed to support modular server, interconnect, and storage components today and throughout the next several years. The enclosure is 10U high and holds up to 16 server and/or storage blades and up to 4 redundant network and storage fabrics utilizing up to 8 interconnect modules. The newest c7000 Platinum Enclosure supports more traffic and users than previous generations, with a 40% increase in bandwidth and twice the storage bandwidth. For example, this blade system has enough bandwidth to support dual 16Gb Fibre Channel connections to each blade server of a fully populated blade chassis without oversubscription. In this test, we installed one Fibre Channel host bus adapter (HBA) mezzanine card. Each blade server can have up to three mezzanine cards installed.

It includes a 7 Tbps high-speed NonStop mid-plane for wire-once connectivity of server blades to network and shared storage. Power is delivered through a pooled-power backplane, and power input flexibility is provided with choices of single-phase AC, 3-Phase AC, and -48V DC.

HP OneView is a recently released software platform with a modern and integrated workspace for converged infrastructure management. It is architected based on how users interact with complex and highly dynamic systems. Tasks and collaboration are automated and streamlined, simplifying the management of both physical and virtual environments.

The HP ProLiant BL460c Gen8 server blade is a dual-socket server that supports the latest generation of Intel Xeon E5-2600 and E5-2600 v2 processors providing 4, 6, 8, 10 and 12 core solutions in each blade server. This blade server has 16 DIMM slots that support 1866 MHz memory and has two x16 PCIe 3.0 I/O expansion slots that support the highest available bandwidth mezzanine option cards, such as 16Gb Fibre Channel and 40Gb Ethernet. The HP ProLiant BL460c Gen8 server blade supports HP SmartMemory, the industry's first Three Rank (3R) 24GB Registered DIMM, which are 25% faster than previous generations.

Brocade 16Gb SAN Switch for HP BladeSystem c-Class

The Brocade 16Gb SAN Switch for HP BladeSystem c-Class provides Gen 5 (16Gb) Fibre Channel support that hot-plugs into the rear of the HP BladeSystem c-Class chassis. It delivers advanced Fibre Channel switch support in a small form-factor, sharing power and cooling provided by the blade chassis.

These embedded Fibre Channel switches provide 16 internal facing ("downlink") ports and 12 external facing "uplink" (SFP+ or SFP) ports per switch. Up to six of these Fibre Channel switches can be installed in the c7000 chassis. These switches support 8Gb and 16Gb BladeSystem Fibre Channel Mezzanine cards.

These FC switches fully integrate into the HP OneView to provide the user with a single management tool to manage all the elements in a BladeSystem infrastructure.

HP MSA 2040 Storage with 16GFC Support

The HP MSA 2040 Storage system is an affordable, 2U storage system that provides shared storage and supports the latest SAN interconnect technologies such as 16Gb and 8Gb Fibre Channel and others.

The HP MSA 2040 can be deployed with a single controller or dual controllers, and supports up to 24 small form factor (2.5-inch) drives or 12 large form factor (3.5-inch) drives that can include SSDs or SAS HDDs. These storage systems provide four host ports per controller as a standard feature, and support 16Gb Fibre Channel host connections.

The MSA 2040 provides several data services features such as snapshots, volume copy functions and remote snap functions.

Test Description and Environment

Demartek ran a read-intensive data warehouse workload with a blade server and storage system. For these tests, we ran the TPC Benchmark™ H (TPC-H), which was used to generate a real-world workload, but not for the purpose of publishing official benchmark results. TPC-H is an industry-standard, decision-support benchmark that simulates broad business intelligence database environments most relevant to information systems that provide organizations with answers to critical business analytics. This test was run with the Gen 5 (16Gb) infrastructure running at 16Gb/sec including the server, mezzanine Fibre Channel HBA, embedded 16Gb Fibre Channel switch and MSA 2040 storage system. The test was repeated for 8Gb Fibre Channel by replacing the 16Gb mezzanine FC HBA inside the blade server with an 8Gb mezzanine FC HBA and by replacing the 16Gb Fibre Channel switch with an 8Gb Fibre Channel switch.

This test was performed by connecting the blade chassis, via the embedded 16GFC switch to two 16Gb Fibre Channel HP MSA 2040 storage systems configured with eight SSDs in each storage system. This provided a solid, low-cost shared storage platform for this test.

Although the diagram on the right shows a fully populated blade chassis, only one server was used for the tests described below.

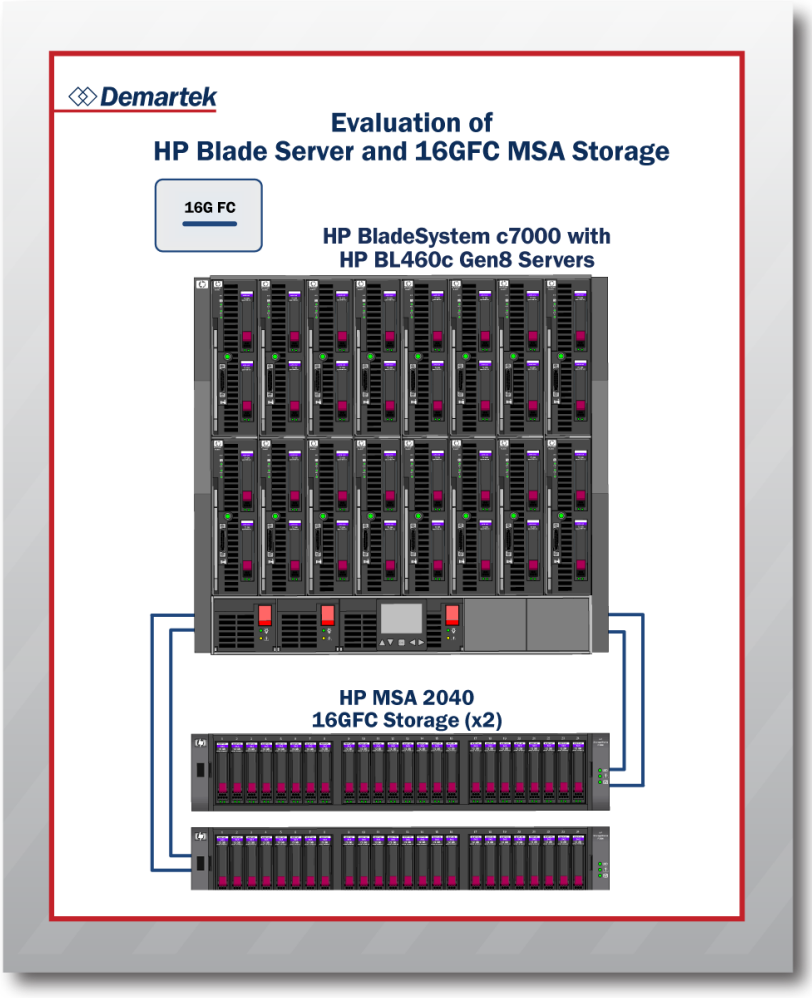

Server

- HP ProLiant BL460c Gen8 Server Blade
- 2x Intel® Xeon® E5-2680, 2.7GHz, 16 total cores, 32 logical processors
- 384GB RAM
- Boot drive: HP 300GB 15K RPM SAS HDD
- Microsoft Windows Server 2012

Fibre Channel HBA

- HP QMH2672 16Gb FC HBA
- HP QMH2572 8Gb FC HBA

Fibre Channel Switch

- 2x Brocade 16Gb SAN Switch for HP BladeSystem c-Class
- 2x Brocade 8Gb SAN Switch for HP BladeSystem c-Class

Storage Array

- 2x HP MSA 2040 Storage
- 16x (8 in each storage system) 400GB 6Gb SAS SSD
  - RAID1 configuration for database volumes
  - RAID1 configuration for log volumes
- 8x (4 in each storage system) 16GFC host ports

Data Warehouse Workload and Performance Results

Database Workload

The read-intensive database workload used consisted of a suite of business-oriented, ad-hoc queries and concurrent data modifications. The queries and the data populating the database were chosen to have broad, industry-wide relevance.

This particular workload consisted of 22 different queries, each exercising a different area of the database. The workload ran the same queries in the same order for each run, with a lower elapsed time for each query indicating a faster system. This benchmark illustrates decision support systems that examine large volumes of data, execute queries with a high degree of complexity, and give answers to critical business questions.

Real vs. Synthetic Workloads

The workload employed in this test used a real database (Microsoft SQL Server) with database tables, indexes, etc., performing actual database transactions. When using real database workloads, I/O rate will vary as the workload progresses because the database performs operations that consume varying amounts of CPU and memory resources in addition to I/O resources. These results more closely resemble a real customer environment.

This is unlike benchmarks that use synthetic workloads that perform the same I/O operations repeatedly, resulting in relatively steady I/O rates which, although potentially faster, do not resemble real customer environments.

Hardware and Software Specifications

The specifications for the software used for this test are listed below. The RAM allocated to the database application was limited to 8Gb in order to force the I/O through the Fibre Channel HBAs and out to the storage system with minimal host memory caching.

Database

- Microsoft SQL Server 2012
- RAM allocated to SQL Server: 8Gb
- Database size: 592 GB
- Log size: 271 GB

Bandwidth Results

For this set of tests, Demartek used a dual-port host connection to the SAN. Multiple runs were performed with six (6) database users.

When running this test with the 16Gb Fibre Channel infrastructure, the time to complete the run was 63% of the time required by the 8Gb infrastructure, or 37% faster. This reduced the time to complete the 6-user workload by approximately four hours.

For many database workloads, time to complete the work is critical because this improvement compounds with more users and more transactions. In large environment, these time reductions can lead to many hours saved.

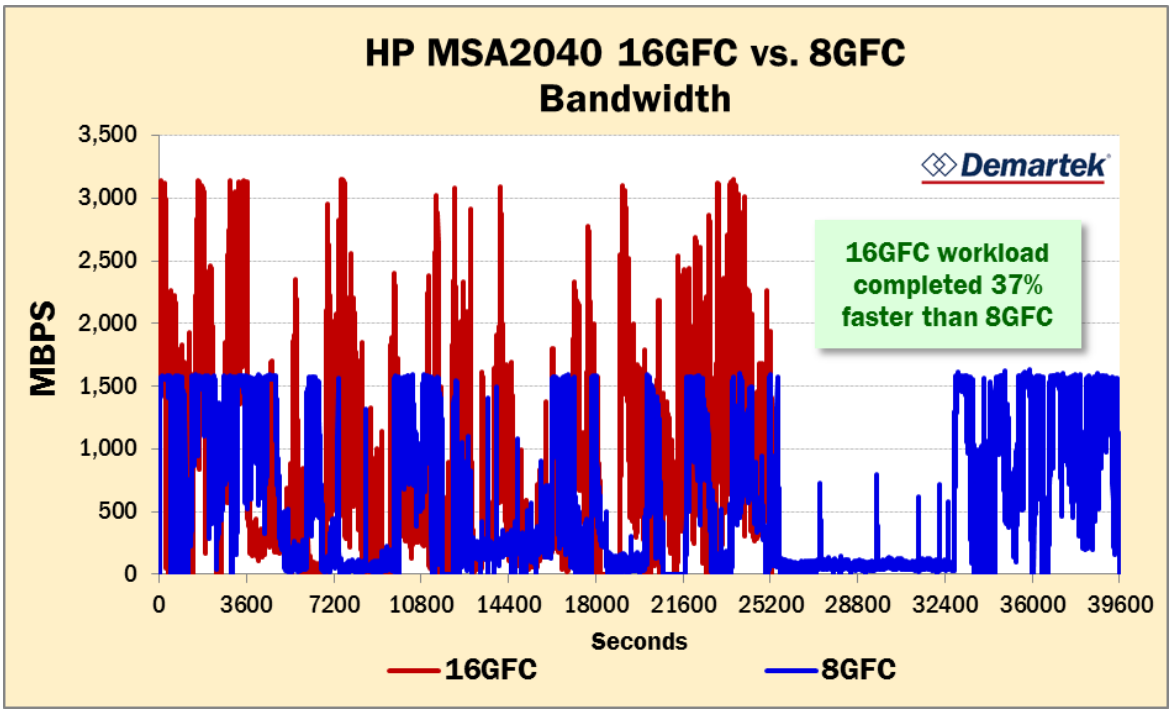

Latency Results

The 16Gb FC infrastructure provided on average approximately a 59% reduction in latency for this workload compared to the 8Gb infrastructure. The average latency for the duration of the tests was 3.68 milliseconds for the 16GFC infrastructure and 8.89 milliseconds for the 8GFC infrastructure. Another way to state this is that the 16Gb infrastructure had less than half of the latency of the 8Gb infrastructure.

Note that the latency pattern is similar between the workloads, with the exception that the workload running with the 16Gb adapter completed the work 37% faster or by approximately four hours.

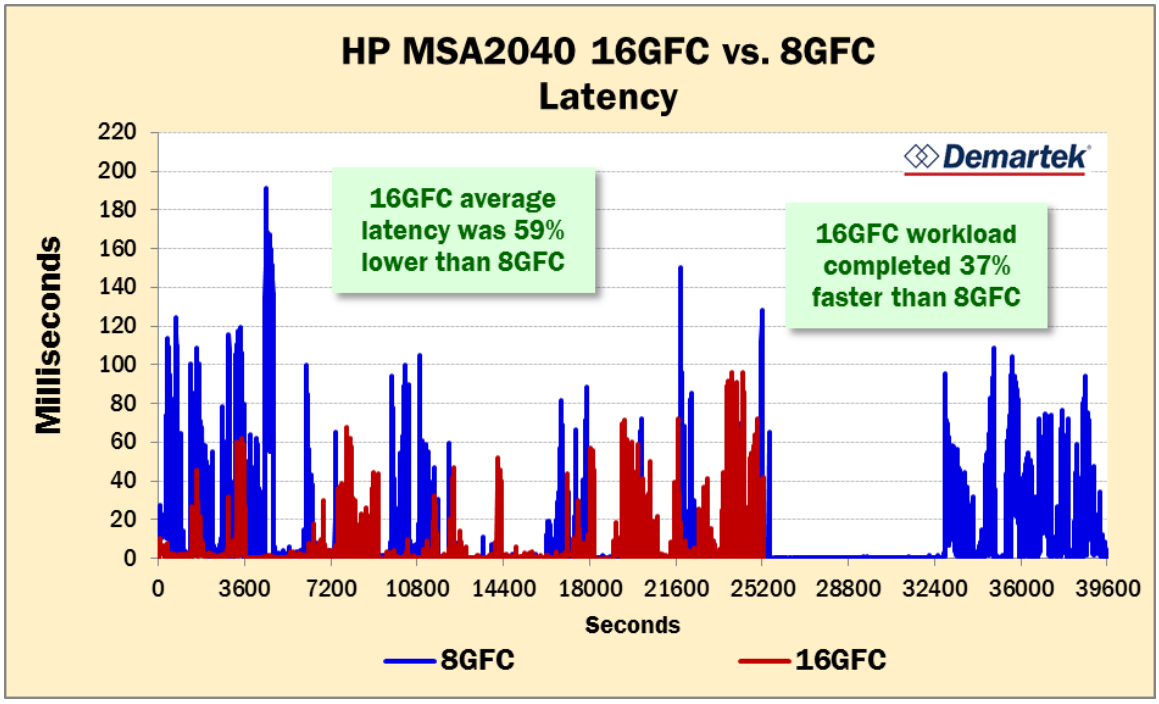

CPU Utilization

The 16Gb infrastructure performed the same work in less time, demonstrated higher bandwidth performance and significantly lower latency. One might ask if this improvement in performance came at the cost of much higher CPU utilization on the host. It turns out that the host CPU did work a bit harder to obtain this additional performance, but not much higher. For the test period, which took 37% less time for the 16Gb infrastructure, the average CPU utilization was 14.3% using the 16Gb infrastructure, and 10.2% using the 8Gb infrastructure, a difference of approximately 4%.

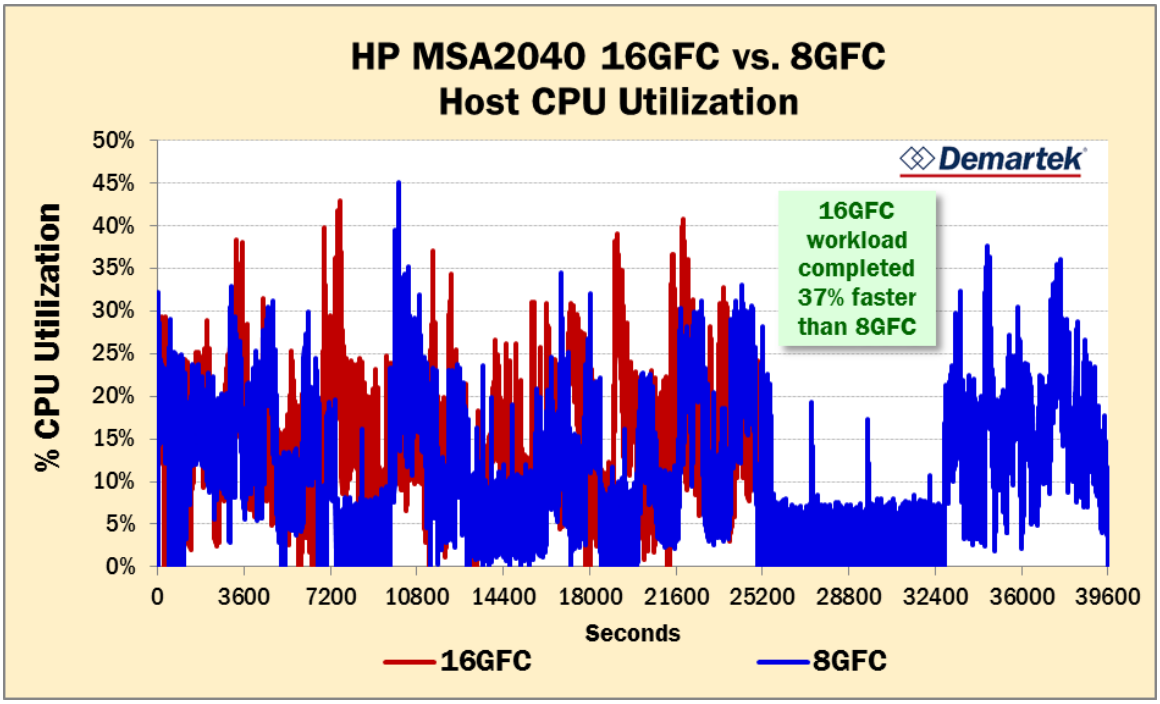

Summary and Conclusion

With the availability of end-to-end 16Gb Fibre Channel infrastructure, IT departments can meet growing performance requirements of virtualized servers, database applications, SSDs, flash caching and other factors.

The new HP BladeSystem c7000 Platinum enclosure integrates next generation technologies to deliver the most advanced architecture with latest end-to-end performance advancements including the new 16Gb Fibre channel switching solutions that were tested in this report. The newest "Platinum" enclosure supports twice the storage bandwidth as previous generations.

The database workload tests using the 16GFC infrastructure achieved higher performance than the same server, storage configuration and workload using 8Gb infrastructure, completing the workload 37% faster and with an average reduction in latency of approximately 59%. The 16GFC infrastructure reduced the time to complete the workload by 4 hours, saving a tremendous amount of time.

16Gb Fibre Channel provides the performance horsepower for both new environments and existing environments that demand higher performance than are available today with older technologies.

This supports HP's claim that "The HP BladeSystem maximizes investment protection by enabling you to easily leverage the latest technologies when they become available by changing just the components that need to be changed instead of the 'rip and replace' approach."

The original version of this document is available at:
http://www.demartek.com/Demartek_HP_BladeServer_MSA_16GFC_Storage_Evaluation_2014-03.html