Lenovo

The Benefits of Emulex 16 Gb Gen6 FC HBAs in Lenovo Solutions

Discusses application requirements for storage

Introduces Emulex 16 Gb Gen6 FC HBAs for Lenovo System x servers

Describes Emulex 16 Gb Gen6 FC HBAs usage scenarios

Provides proof-of-concept performance measurements

Ilya Krutov

LENOVO PRESS

Abstract

Ensuring that business-critical data is available when needed is an ever-growing need in IT. Your systems must store massive amounts of data quickly and retrieve it efficiently. Simultaneously, you must use new technologies that can improve efficiency and take advantage of these technologies within limited budgets. One measure of growing efficiency in recent years is CPU processing power, which far exceeds growth in disk input/output (I/O). For this reason, disk I/O is often the reason for bottlenecks in high-performance applications.

Emulex 16 Gb Gen6 Fibre Channel HBA connectivity combined with other storage technologies, such as Lenovo 16 Gb (Gen 5) FC SAN switches, hybrid storage with solid-state drives and storage tiering, and all-flash storage arrays, can help to balance the performance of the storage subsystem with a server.s processing capabilities.

This Lenovo Press publication describes server performance imbalance that can be found in typical application environments and how to address this issue with the 16 Gb (Gen 6) Fibre Channel host connectivity to provide required levels of performance and availability for storage-intensive applications.

This paper is intended for IT professionals who want to learn more about the Emulex 16 Gb Gen6 FC HBAs for Lenovo System x servers and the benefits that they provide to customers in storage connectivity solutions.

See a list of our most recent publications at the Lenovo Press web site:

http://lenovopress.com

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest, if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Executive summary

Currently, the processor, memory, and I/O subsystem are balanced and not considered to be performance bottlenecks in most systems. The major source of performance issues is related to the storage I/O activity because the speed of traditional storage systems does not match the processing capabilities of the servers. This disparity can lead to a situation where a powerful processor sits idle waiting for the storage I/O requests to complete. This situation wastes the processor's time, which limits user productivity, extends the timeframe for return on investment (ROI), and increases overall total cost of ownership (TCO).

With the virtualization trends in data centers, servers demand higher I/O bandwidth to match the capabilities of multicore processors and large amounts of memory, allowing the higher number of virtual machines (VMs) to be hosted on a single physical system. Higher I/O bandwidth, including storage I/O, can help to achieve better server utilization and a higher VM per server ratio.

Data warehouses and business analytics are additional examples of workloads that require higher storage I/O bandwidth to allow faster data processing. This, in turn, enables making strategic business decisions in a timely manner.

Lenovo® storage solutions that use the Emulex 16 Gb Gen6 FC HBAs can help address these issues by delivering as much as 1.6 million I/O operations per second (IOPS) per port and up to 1.6 GBps throughput per port in each direction. Combined with the reliable high-speed Lenovo 16 Gb (Gen 5) FC storage fabric and solid-state drive technology with storage tiering or all-flash storage arrays, the Emulex 16 Gb Gen6 FC HBAs can help significantly decrease storage I/O response time to match the processing power of the server processors.

Compared to Emulex 8 Gb FC connectivity, the Emulex 16 Gb Gen6 FC HBAs for Lenovo System x® servers can help achieve:

- Up to two times better virtual machine density and a higher number of concurrent users, because of increased storage bandwidth
- Up to 50% fewer servers to support the workload that is specified
- Higher reliability and availability of services, due to fewer components needed to build the solution
- Twice the access speed to business-critical data
- Lower acquisition costs, due to fewer systems and components
- Shorter ROI timeframe and decreased overall TCO, because of more efficient utilization of server resources, as well as lower power, cooling, and management costs

In addition, the Emulex 16 Gb Gen6 FC HBAs can provide up to 5% better performance than the Emulex 16 Gb (Gen 5) FC HBAs, while increasing scalability and lowering power consumption at the same time.

Storage requirements

Choosing the correct storage for application data can be a complex task because you must ensure that critical business and application requirements are met while also optimizing costs. In particular, storage performance capabilities must match the processing capabilities of the server itself to ensure the most efficient utilization of system resources. There is no “one size fits all” approach possible, because different applications have different storage data access patterns.

In general, the factors to consider during the planning process for application data storage include:

- Importance of data (Can I accept the loss of data?)
- Sensitivity of data (Do I need advanced data protection and security?)
- Availability of data (Do I need the data 24 hours per day, 7 days per week?)
- Security of data (Who can read, modify, and delete the data?)
- Data access speed (How quickly do I need to insert and extract the data?)
- Performance or workload capacity (How many IOPS for IOPS-intensive workloads and how many MBps for throughput-intensive workloads do I need?)
- Storage capacity (How much space do I need to store the data?)
- Frequency of access (How often do I need the data?)
- Backup and recovery strategy (How much time do I need to back up and restore the data?)
- Retention policy (How long do I keep the data?)
- Scalability for future growth (Do I expect the workload to increase in the near future?)
- Storage deployment: internal or external (If external, then JBOD or storage system? If storage system, then SAS, iSCSI, FC, or FCoE?)
- Data access pattern (How does the application access the data?):
  - Read- or write-intensive
  - Random or sequential access
  - Large or small I/O requests

Answers to these questions help you to formalize the performance, availability, and capacity requirements for your applications and match these requirements with the appropriate storage design model.

Storage performance

In general, there are two key types of storage applications based on workloads:

- *IOPS-intensive* applications require storage systems to process as many host read and write requests (or I/O requests) per second as possible given the average I/O request size used by this application, which is typically 8 - 16 KBytes. This behavior is most common for OLTP databases.
- *Throughput-intensive* applications require storage systems to transfer to or from the host as many gigabytes of information per second as possible, and they typically use I/O request sizes of 64 KB - 1 MB. These characteristics are commonly inherent to file servers, multimedia streaming, and backup.

Therefore, there are two key performance metrics to evaluate storage system performance: *input/output requests per second (IOPS)* and *throughput (measured in GBps),* depending on application workload.

Table 1 lists typical storage-intensive applications and their workload patterns in a multiuser environment.

Table 1 Typical application workload patterns

Workload Type → ↓ Application Type	Read intensive	Write intensive	IOPS intensive	Throughput intensive	Random access	Sequential access	Good for 16 Gb FC
File server	Yes			Yes	Yes		Yes
OLTP Database	Yes	Yes	Yes		Yes		Yes
Data warehouse	Yes			Yes	Yes		Yes
Email server	Yes	Yes	Yes		Yes		Yes
Medical imaging	Yes			Yes	Yes		Yes
Video on demand	Yes			Yes	Yes		Yes
Streaming media	Yes			Yes		Yes	Yes
Web/Internet	Yes		Yes		Yes		Yes
Web 2.0	Yes	Yes	Yes		Yes		Yes
Backup/archiving		Yes		Yes		Yes	Yes

Compared to 8 Gb Fibre Channel, throughput-intensive applications, such as data warehouses, video on demand, and others, can benefit from 16 Gb Fibre Channel because of the increased bandwidth. Also, IOPS-intensive workloads, such as OLTP, Email, and web applications, can benefit from 16 Gb FC technology because of its higher IOPS capability and lower response time.

Fibre Channel SANs versus converged networks

In the recent past, converged networks were a popular discussion topic because of their ability to carry both LAN and SAN traffic over the same physical infrastructure. Depending on infrastructure requirements, this can help to reduce the number of ports, adapters, and devices that are used, and therefore decrease overall TCO.

However, there is no “one size fits all” approach, and the selection of a storage connectivity approach can be a difficult task. Both dedicated SANs and converged networks have their own strengths and benefits, and the final choice depends on what are the most important requirements in each particular deployment.

Table 2 on page 6 provides an overview of factors that can be considered when deciding on which approach to choose.

Tip: The terms Moderate and High as used in Table 2 are relative indicators for comparison purposes and do not represent any absolute values. For example, values in the TCO row mean that FC SAN typically (but not necessarily always, depending on specific deployment requirements) has a higher TCO than a converged FCoE network.

Table 2 FC SAN versus converged FCoE network

Factor	FC SAN	Converged FCoE network
Storage I/O workload	Medium to heavy	Light to medium
Storage I/O infrastructure	Dedicated	Shared
Traffic isolation	Physical	Logical
Availability	High	Moderate
Performance	High	Moderate
Scalability	High	Moderate
Security	High	Moderate
Storage infrastructure management	Separate	Unified
Acquisition costs	High	Moderate
Management costs	High	Moderate
Operational costs	High	Moderate
TCO	High	Moderate

In general, as shown in Table 2 on page 6, FC SANs provide better performance, availability, scalability, and security, and converged networks can achieve lower TCO.

If the storage workload is light to medium and TCO is the key decision factor, a converged network is an appropriate choice. If the storage workload is moderate to heavy and performance, availability, scalability, and security are the key decision factors, FC SAN is an appropriate choice.

Benefits of 16 Gb Fibre Channel SANs

In general, 16 Gb Fibre Channel means twice the bandwidth of 8 Gb Fibre Channel. As was already described, the processor, memory, and I/O subsystem are balanced and not considered to be performance bottlenecks in the majority of systems. The main source of performance issues tends to be related to storage I/O activity, because the speed of traditional storage systems is significantly lower than the processing capabilities of the servers. Therefore, increasing the speed of storage connections can help to close or minimize this performance gap.

The use of 16 Gb FC connectivity can help to achieve the following benefits:

- Higher performance:
  - Higher IOPS
  - Higher throughput
  - Lower latency
- Infrastructure simplification:
  - Simplified deployment and management
  - Higher VM density/lower number of physical systems
- Improved TCO:
  - Reduced acquisition costs
  - Reduced power and cooling costs
  - Reduced support/maintenance costs

Performance

The maximum theoretical throughput of 16 Gb Fibre Channel is 1.6 GBps, and 8 Gb FC has a throughput of 0.8 GBps. This fact means that throughput-intensive applications can achieve better response time (for example, with OLAP workloads) or support larger numbers of concurrent users on a single server (such as Video on Demand or streaming media workloads).

Also, 16 Gb FC can potentially be beneficial for IOPS-intensive applications: a 16GFC link can deliver up to 200,000 IOPS using the 8 KB I/O blocks, while 8GFC connection can deliver only up to 100,000 IOPS.

Infrastructure simplification

Because the performance of a single 16 Gb FC port is equivalent to the performance of two 8 Gb FC ports, you can reduce the number of ports, cables, and SFP modules, while keeping the same performance levels.

Also, if 8 Gb FC is a limiting factor for overall server performance, customers can support heavier workloads on a same system by moving to 16 Gb FC; thereby increasing VM density or concurrent user support while reducing the overall number of physical systems.

Improved TCO

With the lower number of systems and components that are used to build the infrastructure, customers can reduce acquisition costs and operational, management, and support costs due to a simplified infrastructure.

This paper describes several possible deployment scenarios where Lenovo 16 Gb FC infrastructure can provide certain benefits:

- Online transaction processing
- Decision support
- Video on Demand
- Mixed virtualized workloads

For more information, see “Deployment scenarios of 16 Gb (Gen 6) FC host connectivity” on page 11.

Emulex 16 Gb Gen6 FC HBAs for Lenovo System x servers

The Emulex 16 Gb (Generation 6) Fibre Channel (FC) host bus adapters (HBAs) are an ideal solution for all Lenovo System x servers requiring high-speed data transfer in storage connectivity for virtualized environments, data backup, and mission-critical applications. They are designed to meet the needs of modern networked storage systems that utilize high performance and low-latency solid-state drives for caching and persistent storage as well as hard disk drive arrays.

The Emulex 16 Gb Gen6 FC HBAs feature ExpressLane, which prioritizes mission-critical traffic in congested networks to ensure maximum application performance on flash storage arrays. These HBAs also seamlessly support Brocade ClearLink diagnostics through Emulex OneCommand Manager, ensuring the reliability and management of a storage network, when connected to Brocade Gen 5 FC SAN fabrics.

The Emulex 16 Gb Gen6 FC HBAs are shown in Figure 1.

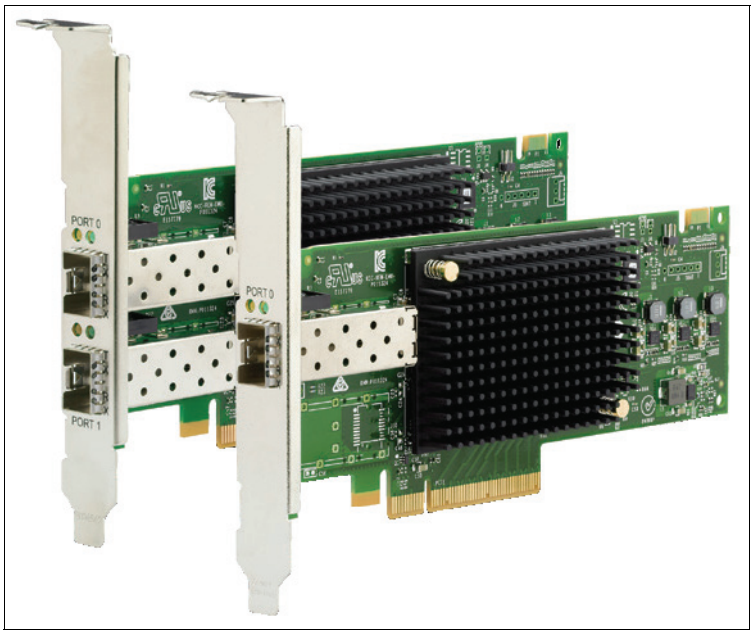

Figure 1 Emulex 16 Gb Gen6 FC Single-port (right) and Dual-port (left) HBAs (without SFP+)

The Emulex 16 Gb Gen6 FC HBAs have a 32 Gbps FC ASIC, which can achieve 1.6M IOPS on a single port by using Emulex's Dynamic Multicore architecture, which dynamically scales HBA resources to any port that needs it. This is essential when ports are used in active-standby mode.

The Emulex 16 Gb Gen6 FC HBAs can provide near limitless scalability to support maximum VM density, with 2x more on-chip resources and bandwidth. These low latency HBAs can also improve your VDI experience, providing noticeable improvements during boot storms, and allow faster data warehousing and meet the massive bandwidth requirements of flash storage arrays.

Table 3 lists the ordering information for the Emulex 16 Gb Gen6 FC HBAs.

Table 3 Ordering information for the Emulex 16 Gb Gen6 FC HBAs

Description	Part number	Feature code
Emulex 16Gb Gen6 FC Single-port HBA	01CV830	ATZU
Emulex 16Gb Gen6 FC Dual-port HBA	01CV840	ATZV

The part numbers for the Emulex 16 Gb FC HBAs include the following items:

- An FC HBA adapter with one or two 16 Gb (16/8/4 Gbps speeds) FC SW SFP+ installed
- 3U (standard) and 2U (low-profile) adapter brackets
- Publications package

The Emulex 16 Gb Gen6 FC HBAs have the following features:

- Maximum performance with up to 1.6 million input/output operations per second (IOPS) to support larger server virtualization deployments and scalable cloud initiatives, and performance to match new multicore processors, SSDs/flash storage, and faster server host bus architectures.

- Supports Brocade Clearlink diagnostics, which helps ensure optical and signal integrity for Fibre Channel cables and optics by validating the health, reliability and performance of the network prior to, and after, deployment. Allows the IT administrator to detect faulty cables and optics in minutes versus hours. Brocade ClearLink is also seamlessly integrated into Emulex OneCommand.
- Offer end-to-end Quality of Service (QoS) application prioritization with ExpressLane technology, which allows customers to prioritize faster storage traffic (such as SSDs) ahead of slower traffic (such as spinning hard drives), which alleviates potential bottlenecks from slow storage.
- Frame-level multiplexing and out-of-order frame reassembly increases link efficiency and maximizes HBA performance.
- vScale performance and scalability: Multicore ASIC engine with eight cores supports 255 VFs, 1024 MSI-X, and 16127 logins/open exchanges for maximum VM density.
- The Emulex OneCommand Manager enterprise class management application features a multiprotocol and cross-platform architecture that provides centralized management of all Emulex HBAs. VMware vCenter plug-in provides OneCommand support within a VMware environment.
- GreenState power efficiency reduces data center power consumption and associated operational expenses by delivering exceptional power to port ratios.
- End-to-end data protection with hardware parity, CRC, ECC, and other advanced error checking and correcting algorithms, which ensures that data is safe from corruption.
- Support forward error correction (FEC), a new Gen 6 standard feature that provides enhanced data reliability and performance by automatically detecting and recovering from bit errors.
- T10-PI data integrity with high-performance offload provides end-to-end data corruption protection.
- Rock-solid reliability and thermal characteristics, which are essential for mission-critical, cloud, and virtualized applications.
- Emulex HBAs are renowned for reliability, ensuring maximum SAN uptime. Their "it just works" reputation is based on 17 million installed ports with proven industry-leading reliability of 10 million hours field Mean Time Between Failures (MTBF).
- Support for Message Signaled Interrupts eXtended (MSI-X) improves host utilization and enhances application performance.
- Support for 16 Gb, 8 Gb, and 4 Gb FC devices.
- Comprehensive virtualization capabilities with support for N_Port ID Virtualization (NPIV).
- A common driver model allows a single driver to support all Emulex HBAs on a given OS.
- Reduces the number of cards, cables, and PCIe slots required.
- Exceptional performance per watt and price/performance ratios.
- Integrates seamlessly into existing SANs.
- Allows application of SAN best practices, tools, and processes with virtual server deployments.
- Ensures data availability and data integrity.
- Universal boot capability allows the appropriate boot environment to be automatically selected for any given hardware.
- Boot from SAN capability reduces systems management costs and increases uptime.

- Detailed and real-time event logging and tracing enables quick diagnosis of SAN problems.
- The beaconing feature flashes the HBA LEDs, simplifying their identification within server racks.
- The environmental monitoring feature helps optimize SAN availability.

Table 4 compares features of Emulex 16 Gb (Gen 6), 16 Gb (Gen 5), and 8 Gb FC HBAs.

Table 4 Emulex 16 Gb (Gen 6), 16 Gb (Gen 5), and 8 Gb FC HBAs feature comparison

Feature	16 Gb FC (Gen 6)	16 Gb FC (Gen 5)	8 Gb FC
Part numbers	01CV830 01CV840	81Y1655 81Y1662	42D0485 42D0494
Host interface	PCIe 3.0 x8	PCIe 3.0 x8	PCIe 2.0 x8
IOPS performance	1.6 M IOPS	1.2 M IOPS	0.2 M IOPS (per port)
32 Gb FC SFP+ transceiver support	Yes	No	No
16 Gb FC SFP+ transceiver support	Yes	Yes	No
8 Gb FC SFP+ transceiver support	No	Yes	Yes
32GFC speed support	Yes[a]	No	No
16GFC speed support	Yes	Yes	No
8GFC speed support	Yes	Yes	Yes
4GFC speed support	Yes	Yes	Yes
ClearLink support	Yes	Yes	No
ExpressLane support	Yes	Yes	No
Logins and Exchanges	16,127	8,192	4,096
SR-IOV support	255 VFs	255 VFs	No

a. Requires a 32 Gb FC SFP+ transceiver.

Emulex 16 Gb Gen6 FC HBAs offer better performance and scalability with lower power consumption, while fitting into the same price band, than Emulex 16 Gb (Gen 5) FC HBAs.

The Emulex 16 Gb Gen6 FC HBAs have the following specifications:

- I/O controller: Emulex Engine 501 (XE501) I/O Controller (IOC)
- Host interface: PCIe 3.0 x8
- Ports: Single-port and dual-port SFP+ based adapters
- Link speed: Support for 16GFC, 8GFC and 4GFC link speeds, which are automatically negotiated
- Data rate: 14.025 Gbps (1600 MBps), 8.5 Gbps (800 MBps), and 4.25 Gbps (400 MBps) autosensing (per port), with full duplex
- Performance: Up to 1,600,000 IOPS

- Industry standards:
  - Current ANSI/IETF standards: FC-PI-4, FC-PI-5, FC-PI-6, FC-FS-3, FC-LS-2, FC-GS-6, FC-DA, FC-DA2, FCP-4, SPC-4, SBC-3, and SSC-4
  - Legacy ANSI/IETF standards: FC-PH, FC-PH-2, FC-PH-3, FC-PI, FC-PI-2, FC-PI-3, FC-FS, FC-GS-2/3/4/5, FCP-2/3, FC-HBA, FC-TAPE, FC-MI, SPC-3, SBC-2, SSC-2, and SSC-3
- Topology: Point-to-point and switched fabric
- Supported media: Hot-pluggable 16 Gbps Fibre Channel SFP+ short wave optical transceivers (850 nm) with LC connectors (included with the adapters)
- Distance support:
  - Operating at 16 Gbps:
    - Up to 15 m on 62.5/125 µm OM1 Multi-Mode Fiber (MMF)
    - Up to 35 m on 50/125 µm OM2 MMF
    - Up to 100 m on 50/125 µm OM3 MMF
    - Up to 125 m on 50/125 µm OM4 MMF
  - Operating at 8 Gbps:
    - Up to 21 m on 62.5/125 µm OM1 MMF
    - Up to 50 m on 50/125 µm OM2 MMF
    - Up to 150 m on 50/125 µm OM3 MMF
  - Operating at 4 Gbps:
    - Up to 70 m on 62.5/125 µm OM1 MMF
    - Up to 150 m on 50/125 µm OM2 MMF
- Management software:
  - Emulex AutoPilot Installer automates the HBA installation process and reduces time to deployment and administrative costs. Automated installation and configuration of driver and management tools simplifies deployment of multiple adapters within Windows environments. A single installation of driver and management tool eliminates multiple reboots and ensures that each component is installed correctly and ready to use.
  - The Emulex OneCommand Manager application enables centralized discovery, monitoring, reporting, and administration of Emulex HBAs and CNAs on local and remote hosts. Powerful automation capabilities facilitate remote driver parameter, firmware, and boot code upgrades. In addition to the GUI interface, management functions can also be performed through a scriptable command-line interface (CLI) and a web browser.
  - Emulex management instrumentation complies with Open Management Standards, such as SMI-S and common HBA API support, which enables seamless upward integration into enterprise storage and server management solutions.

Deployment scenarios of 16 Gb (Gen 6) FC host connectivity

Even if the storage system does not have 16 Gb Fibre Channel ports, in certain cases 16 Gb FC host connectivity can provide significant advantages over the 8 Gb FC deployments. These advantages are:

- Lower acquisition costs
- Simplified cabling
- Lower operational costs
- Lower maintenance and support costs

For the purpose of validating the benefits of the Emulex 16 Gb Gen6 FC HBAs, we set up a lab environment consisting of the following components:

- Lenovo System x3650 M5 server with three Emulex adapters installed:
  - Dual-port 8 Gb FC HBA
  - Single-port 16 Gb (Gen 5) FC HBA
  - Dual-port 16 Gb (Gen 6) FC HBA
- Lenovo B Series 16 Gb (Gen 5) FC SAN switch
- Storwize V7000 for Lenovo with 8 Gb FC ports

We used four 8 Gb FC ports (two ports on each node canister) on the Storwize V7000 to connect it to the 16 Gb (Gen 5) FC SAN switch, and we also connected the FC HBAs to the same switch, as shown in Figure 2.

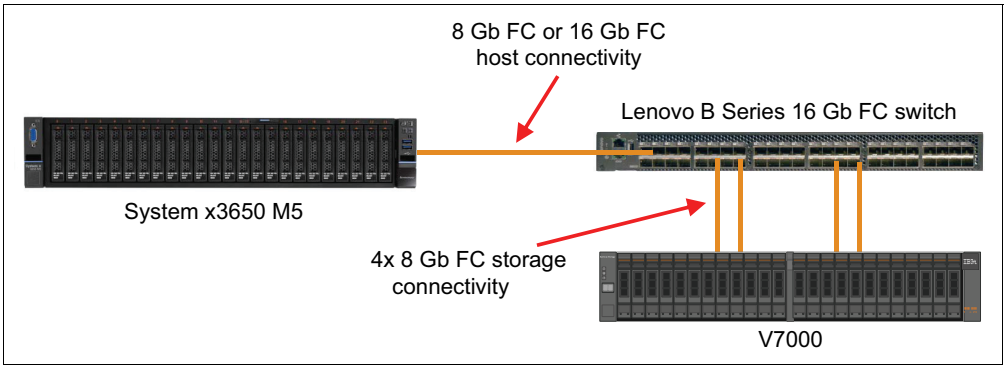

Figure 2 test environment

The following subsections provide suggestions for where to deploy 16 Gb (Gen 6) FC host connectivity to help utilize the potential benefits of 16GFC speeds compared to 8GFC speeds. The following scenarios are described:

- "OLTP databases"
- "Decision support" on page 13
- "Video on Demand" on page 15
- "Mixed workloads" on page 15

OLTP databases

Online transaction processing (OLTP) is a multi-user, memory-, processor-, and storage I/O-intensive, random workload. It is characterized by a large number of small read and write storage I/O requests (typically four or eight kilobytes and 70/30 read/write ratio) generated by transactions originated by multiple users. The transactions are relatively simple; however, every single transaction can generate dozens of physical storage I/O requests depending on the transaction type, application architecture, and business model used.

The key performance indicator of transactional systems is response time: the customers expect to get the requested product information or place an order quickly. If the expectations are not met then the chance that the customer will go to a competitor is high. Because of that, storage I/O performance is considered an important factor to ensure that response time goals are met, and to keep other system resources (CPU and memory) at high utilization and not waiting for the data.

We validated OLTP storage performance for various fabrics by using 8 KB block random 70% read / 30% write operations. We measured the response time on a moderately loaded storage system, as shown in Figure 3 on page 13, and then maximized the load to reach the maximum storage performance capacity for the FC fabric (see Figure 4 on page 13).

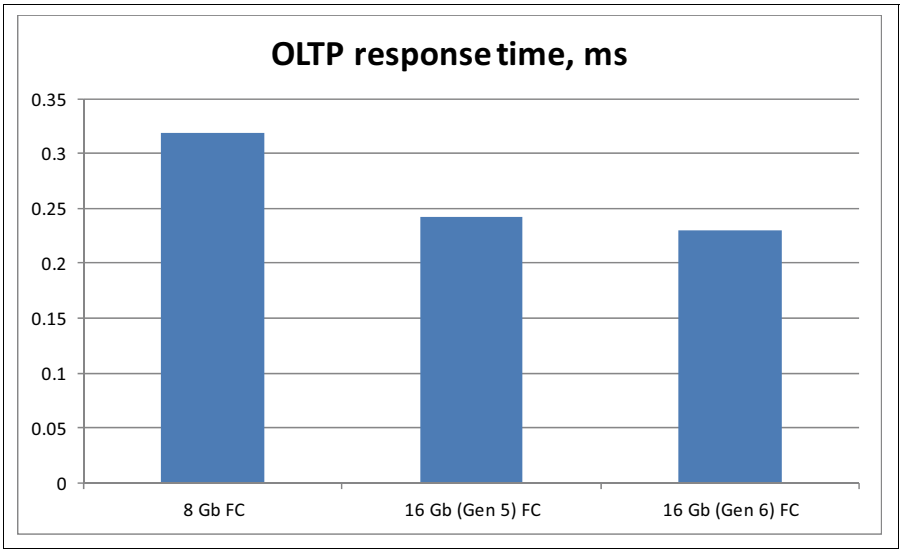

Figure 3 OLTP response time under a moderate workload

Based on our measurements, the Emulex 16 Gb Gen6 FC HBAs delivered 28% lower response time than 8 Gb FC HBAs. Also, the Emulex 16 Gb Gen 6 FC HBAs delivered almost 5% better response time than the Emulex 16 Gb (Gen 5) FC HBAs.

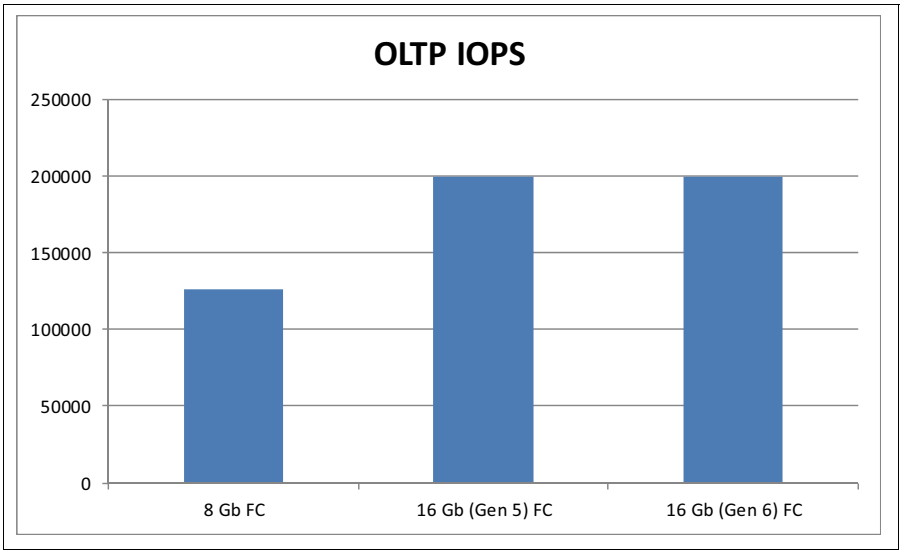

Figure 4 OLTP storage performance capacity

When the storage fabric connectivity became a limiting factor for performance scaling, the Emulex 16 Gb Gen6 FC HBAs achieved over 50% higher IOPS with 35% lower response time, compared to 8 Gb FC HBAs.

For OLTP workloads, Emulex 16 Gb Gen6 FC HBAs can help achieve better utilization of the potential of a storage system by removing the storage bottleneck caused by 8GFC speeds.

Decision support

Data warehouses are commonly used with online analytical processing (OLAP) workloads in decision support systems; for example, financial analysis. Unlike OLTP, where transactions are typically relatively simple and deal with small amounts of data, OLAP queries are much

more complex and process large volumes of data. By its nature, the OLAP workload is sequential read-intensive and throughput-intensive, while it becomes randomized in a multi-user environment.

OLAP databases are normally separated from OLTP databases, and OLAP databases consolidate historical and reference information from multiple sources. Queries are submitted to OLAP databases to analyze consolidated data from different points of view to make better business decisions in a timely manner.

For OLAP workloads, it is critical to have a fast response time to ensure that business decisions support an organization's strategy and are made in a timely manner in response to changing market conditions. Delays might significantly increase business and financial risks; therefore, storage I/O capabilities must match the performance of other server subsystems to ensure that queries are processed as quickly as possible.

In our testing scenario, we simulated decision support queries by using 1 MB block random read operations and assumed ten I/O requests per query on average. We measured the number of queries per second and the average query time. The results are shown in Figure 5.

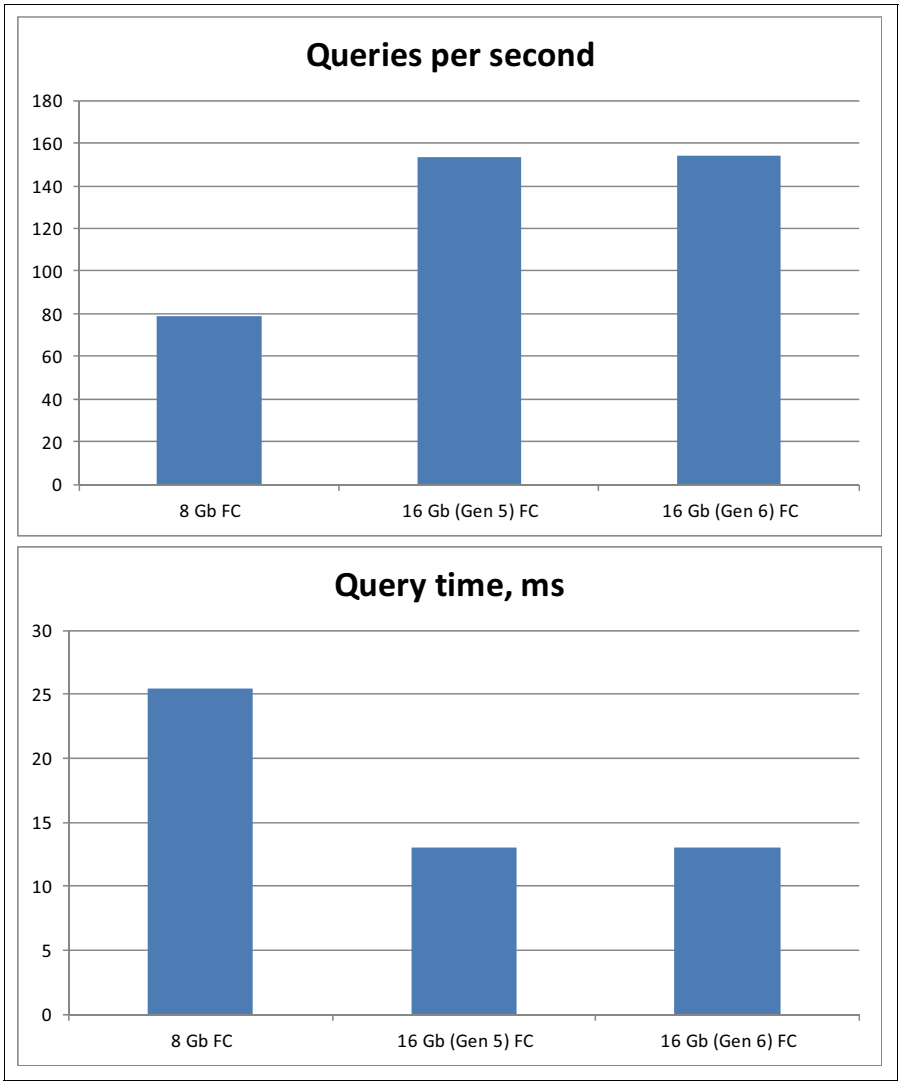

Figure 5 Decision support workload

The results prove that the Emulex 16 Gb Gen6 FC HBAs, as well as Emulex 16 Gb (Gen 5) FC HBAs, are capable of supporting up to almost two times more OLAP queries per second with up to 50% less processing time compared to Emulex 8 Gb FC fabric.

Higher-speed technology provides clear benefits: queries can be completed much faster compared to lower-speed technologies. Also, server utilization increases by processing more queries per second, helping to get faster return on investments.

Video on Demand

Video on Demand (VOD) is traditionally a sequential throughput-intensive workload. However, in a multiuser environment – where every user receives their own data stream watching different content or even the same content with some delay (for example, a recently published new movie) – the workload becomes randomized, which requires faster response time, along with sufficient bandwidth, to ensure a better user experience and smoother video playback. In general, video streaming applications use 64 KB (or larger) I/O blocks to interact with the storage system.

In our test environment, we simulated a VOD workload with 512 KB random read operations, and observed the results that are shown in Figure 6.

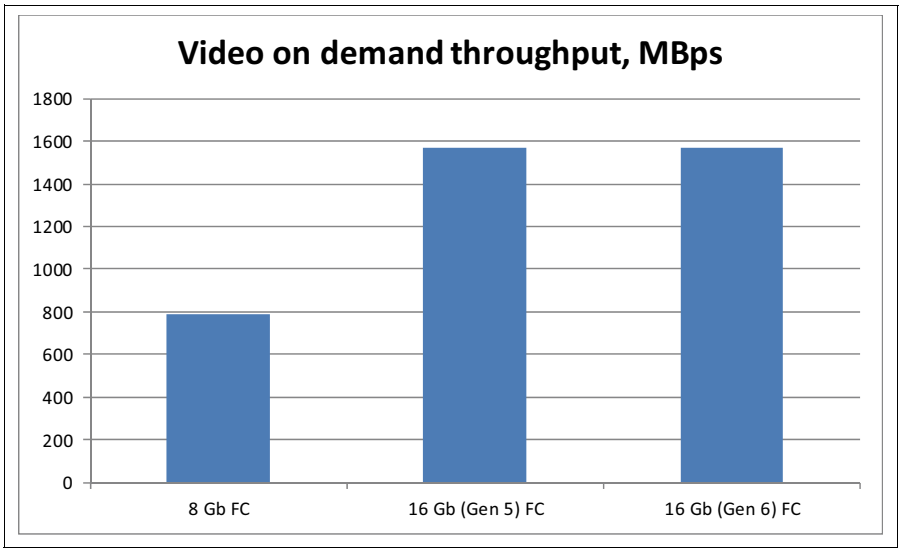

Figure 6 Video on Demand workload

Clearly, the higher bandwidth provided by the 16 Gb FC fabric allows up to 100% more throughput, which helps achieve up to 100% more user streams per server, provided other server subsystems have sufficient user handling capacity and do not represent a bottleneck.

Mixed workloads

In a virtualized environment, we can observe different kinds of workloads and storage access patterns, even on the same physical server. For the purpose of our validation, we defined four types of workloads:

- Database server (random 70% read / 30% write operations using 8 KB blocks)
- E-mail server (random 50% read / 50% write operations using 32 KB blocks)
- Web server (random read operations using 16 KB blocks)
- File server (random and sequential 75% read / 25% write operations using 64 KB blocks)

The results of our validation are shown in Figure 7.

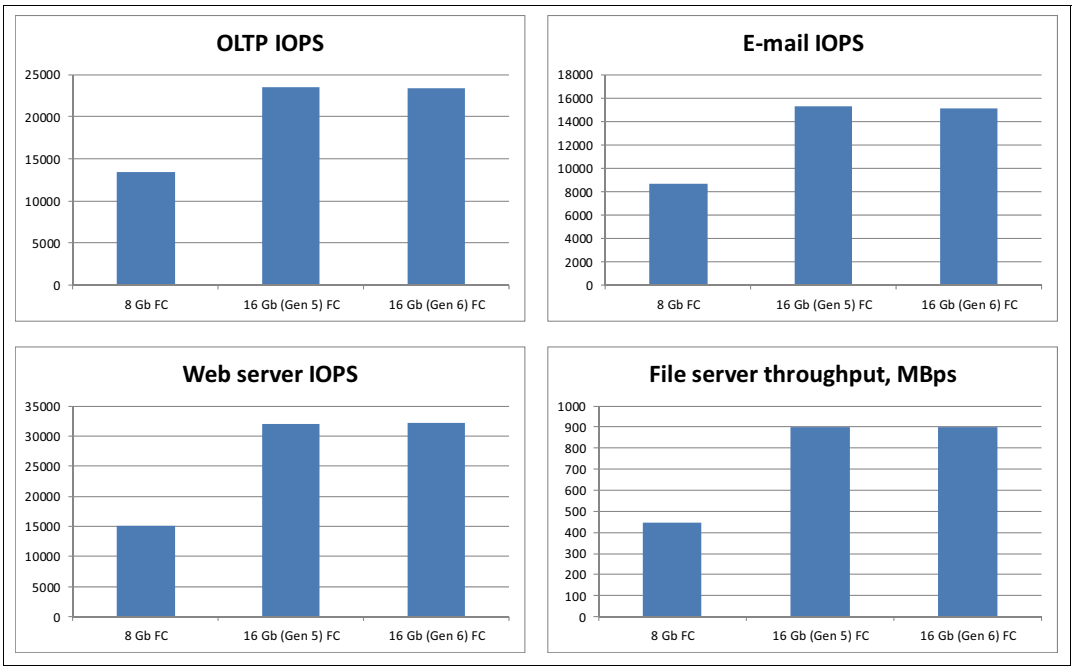

Figure 7 Mixed workloads - performance capacity

For OLTP and e-mail applications, both 16 Gb Gen6 FC and 16 Gb (Gen 5) FC HBAs provide up to 75% more performance capacity, while web server and file server workloads doubled the performance, compared to 8 Gb FC fabric.

The results show that 16 Gb FC connectivity provides an opportunity to host up to two times more virtual machines compared to 8 Gb FC fabric. In addition, we have observed a dramatic decrease in response time when moving to 16 Gb FC fabric, as shown in Figure 8.

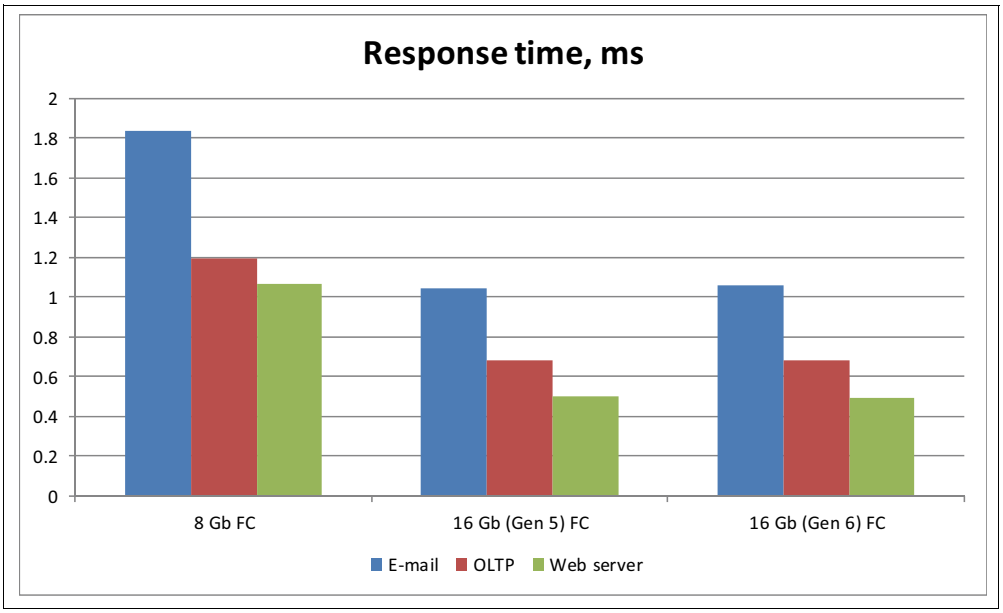

Figure 8 Mixed workloads - response time

Conclusion

Storage I/O represents one of the most common bottlenecks in many high-performance applications. 16 Gb (Gen 6) FC host connectivity can provide up to 1.6 GBps per port for throughput-intensive workloads and up to 1.6 million IOPS per port for IOPS-intensive workloads, allowing customers to potentially eliminate storage I/O bottlenecks in a system.

With 16 Gb (Gen 6) FC host connectivity combined with the 16 Gb (Gen 5) FC storage fabric, customers can deploy fewer servers with better server utilization, including twice the density of virtual machines and a higher number of concurrent users, compared to 8 Gb FC solutions.

Compared to Emulex 8 Gb FC, Emulex 16 Gb Gen6 FC technology can help achieve:

- Up to two times better virtual machine density and more concurrent users, due to increased storage bandwidth.
- Up to 50% fewer servers needed to support the same workload.
- Higher reliability and availability of the services due to fewer solution components.
- Twice faster access to the business critical data.
- Lower acquisition costs due to fewer systems and components needed.
- Shorter ROI time frame and decreased overall TCO with the efficient utilization of server resources and lower power, cooling, and management costs.

In addition, the Emulex 16 Gb Gen6 FC HBAs can provide up to 5% better performance than the Emulex 16 Gb (Gen 5) FC HBAs, while increasing scalability and lowering power consumption at the same time.

Authors

This paper was produced by the following team of specialists:

Ilya Krutov is a WW DCG Sales Enablement Program Manager with Lenovo. He manages and produces presale and post-sale technical publications on various IT topics, including Lenovo x86 rack and blade servers, server operating systems, virtualization and cloud, networking, storage, and systems management. Ilya has 20 years of experience in the IT industry, backed by professional certifications from Cisco Systems, IBM, and Microsoft. During his career, Ilya has held a variety of technical and leadership positions in education, consulting, services, technical sales, marketing, channel business, and programming. He has written more than 200 books, papers, and other technical documents. Ilya has a Specialist's degree with honors in Computer Engineering from the Moscow State Engineering and Physics Institute (Technical University).

Thanks to the following people for their contributions to this project:

From Lenovo:

- Manzar Alam
- Mark T. Chapman
- Tony Meli
- Dean Price
- David Watts

From Broadcom:

- Jeff Lin

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on November 10, 2016.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp0515

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®	Lenovo(logo)®	System x®

The following terms are trademarks of other companies:

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.