



**Hewlett Packard
Enterprise**

Benefits and application of 16 Gbps Fibre Channel

Accelerate your approach to data transfer

Contents

The need for quicker data transfer at minimal costs.....	2
Making the decision to enable a converged data exchange center.....	3
Key benefits of 16 Gbps FC speed	3
Reduced number of links.....	3
Reduced power consumption and TCO	4
The need for speed.....	5
Increased number and size of applications	5
Virtualized servers	5
Multi-core processors.....	6
Larger memories in servers.....	6
Solid state disks.....	6
Faster PCIe	6
Storage systems with 16 Gbps FC interfaces.....	7
HPE Smart SAN—target-driven peer zoning for HPE 3PAR StoreServ Storage implementations.....	7
Aggregation of traffic.....	7
Applications of 16 Gbps FC.....	8
High-bandwidth applications.....	8
Data migration	8
Virtual Desktop Infrastructure	9
Solid state disks.....	9
16 Gbps FC vs. FCoE	9
Demartek key findings.....	9
We recommend	9
Top four reasons to partner with HPE	10
1. Valuable experience.....	10
2. Innovation at every step	10
3. Seamless integration	10
4. Proven success.....	10

Storage area networks (SANs) are constantly in the need for fail-proof, high-performance centers where a large amount of data exchange is the only constant. Several technology advances are pushing up SAN bandwidth demands, including application growth, server virtualization, multi-core processors, PCI Express 3.0, increased memory, and solid state disks (SSDs). The amount of information that is managed, shared, and protected continues to grow. The industry is waking up to the importance of virtualizing data centers and bringing together a converged infrastructure to orchestrate their IT setups. The need for a quicker data exchange has never been more necessary than right now. This paper discusses these industry-transforming dynamics from an industry perspective, and explains how new interfaces allow for faster technology and quicker data transfers at minimal costs.

The need for quicker data transfer at minimal costs

For a converged and virtualized environment to exist, there is immense pressure on the IT organization to manage fewer devices, utilize less energy and resources, and transfer data faster, all while handling similar or higher volumes of work than a traditional environment. Storage array migration, disaster recovery, Virtual Desktop Infrastructure (VDI), and Inter-Switch Links (ISLs) are just some of the typical applications that demand higher bandwidth. When large chunks of data need to be transferred between data centers, arrays, or sites, a faster link can allow the same tasks to be completed in less time.

For many organizations, the way forward is to look at interfaces that will enable speedy data transfer and throughput. The new 16 Gbps Fibre Channel (FC) is instrumental in revolutionizing data throughput. Technology that enables faster data transfer in a virtualized environment bring about results that are easy to see, as the same set of tasks can be accomplished with fewer links, less devices need to be managed, and less power is consumed.

Table 1. FC speed characteristics

Speed	Throughput (Mbps)	Line rate (Gbps)	Encoding	Retimers in the module	Transmitter training
1 GFC	100	1.0625	8b/10b	No	No
2 GFC	200	2.125	8b/10b	No	No
4 GFC	400	4.25	8b/10b	No	No
8 GFC	800	8.5	8b/10b	No	No
10 GFC	1200	10.53	64b/66b	Yes	No
16 GFC	1600	14.025	64b/66b	Yes	Yes

Note

The products and configurations as described in this paper relate to the industry in-general and do not necessarily represent or indicate support by HPE unless explicitly mentioned.

Making the decision to enable a converged data exchange center

Offering considerable improvements from previous FC speeds, 16 Gbps FC uses 64b/66b encoding, retimers in modules, and transmitter training. Doubling the throughput of 8 Gbps to 1,600 Mbps, it uses 64b/66b encoding to increase the efficiency of the link. 16 Gbps FC links also use retimers in the optical modules to improve link performance characteristics, and electronic dispersion compensation and transmitter training to improve backplane links. The combination of these technologies enables the 16 Gbps FC to provide some of the highest throughput density in the industry, making data transfers smoother, quicker, and cost-efficient.

The 16 Gbps FC multimode links are designed to meet distance requirements for a majority of data centers. With the standardization of the OM4 fiber, Fibre Channel has standardized the supported link distances over OM4 fiber and 16 Gbps FC to support 125 meters. If a 16 Gbps FC link needs to go farther than these distances, a single-mode link can be used that supports distances up to 10 kilometers. A wide range of supported link distances enables the 16 Gbps FC to mold right into any data environment as the business demands.

Table 2. Link distance with speed and fiber type (meters)

Speed	OM1 link distance 62.5 μm core 200 MHz*km	OM2 link distance 50 μm core 500 MHz*km	OM3 link distance 50 μm core 2000 MHz*km	OM4 link distance 50 μm core 4700 MHz*km	OS1 link distance 9 μm core ~infinite MHz*km
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
10 GFC	33	82	300	*	10,000
16 GFC	15	35	100	125	10,000

Note

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

Key benefits of 16 Gbps FC speed

Speed allows for more output in a fraction of the time, paving the way for an easy transition to saving costs and increasing efficiency. By doubling up on speed, 16 Gbps FC reduces the time it takes to transfer data between two ports. When more work is accomplished by a single server or storage device, fewer servers, HBAs, links, and switches are needed to accomplish the same task. A typical data center, whether converged or otherwise, with the use of 16 Gbps FC, would consume less power per bit and provide easier cable management.

Reduced number of links

As with other FC speeds, the first application of new speeds is on ISLs between the switches. Large fabrics are composed of many switches that are connected via multiple ISLs. Reduction of the number of ISLs between switches is a benefit of the higher speed. HPE B-series switches support trunking of up to 8 links of 16 Gbps FC to yield a 128 Gbps FC link between any two switches. These trunks can grow from 16 Gbps FC to 128 Gbps FC in 16-Gigabit increments.

The high-speed links of 16 Gbps FC eliminate tens or hundreds of ports as compared to the 8 Gbps FC fabric. Savings take place when the number of HBAs, switches, and end devices gradually decrease with 16 Gbps FC.

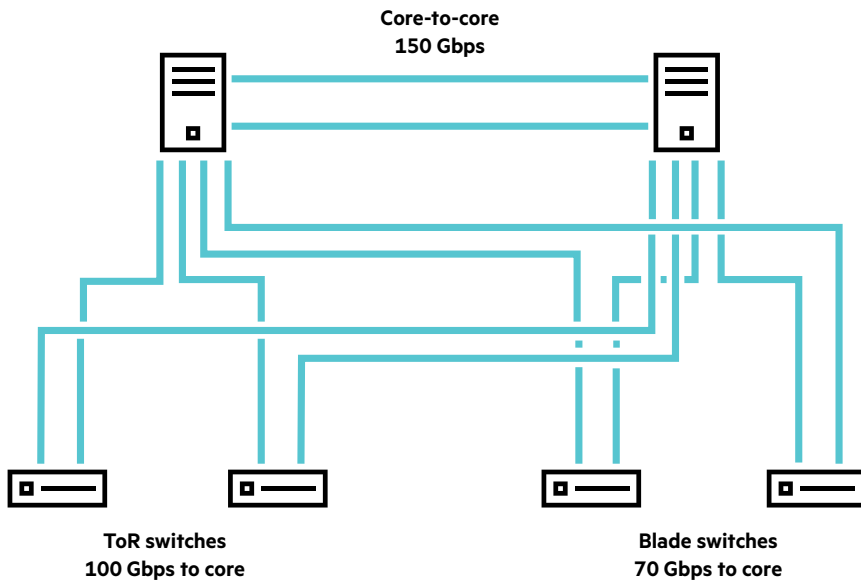


Figure 1. Network design implications

	8 Gbps FC links	16 Gbps FC links
ISLs from ToR switch to core	16	8
ISL from blade switch to core	10	5
Core to core	24	12
Total ISLs	50	25
Total ports	100	50

Reduced power consumption and TCO

Besides allowing for a reduction in equipment that dramatically cuts power consumption, 16 Gbps FC also reduces the power required to transfer bits on the link. When the cost of cabling and operating expenses (OPEX) such as electricity and cooling are considered, the TCO is often less when links are run at twice the speed. The goal of 16 Gbps FC designs is for a 16 Gbps FC port to consume less power than 2 x 8 Gbps FC links that deliver the same throughput. Initial estimates for power consumption show the 16 Gbps FC SFP+ consuming 0.75 watts of power while the 8G FC SFP+ consuming 0.5 watts of power. These estimates show that a 16 Gbps FC link consumes 25 percent less power than 2 x 8 Gbps FC ports.

Easier cable management

Cable management becomes a much simpler process if fewer links are needed. Managing hundreds of cables from a single switch or bundles of cable is not an easy task. A virtualized center spells less dependency on hardware. 16 Gbps FC enables a reduction of cables, avoiding the tedious tasks of troubleshooting and re-cabling. The cost of cabling is significant and users can pay over \$300 USD per port in structured cabling environments. Reducing the number of links by using fast 16 Gbps FC links brings down the cost and work involved in cable management.

The HPE B-series 16 Gb Directors are supported by incredibly dense solutions for cable management. The HPE B-series 16 Gb Directors are supported by incredibly dense solutions for cable management. To reduce the bulk of cables, Corning offers MTP to LC fan out cables, which contain 12 fibers with an MTP connector at one end and 6 duplex (paired) LC connectors at the other. The MTP to LC cable harnesses reduce cable bulk while allowing direct connection into standard equipment.

The need for speed

The principal benefit of 16 Gbps FC is increased performance at double the data rate of 8 Gbps FC and 40 percent faster than 10GbE. This feature allows better functioning of multiple technology advances including:

- Increased number and size of applications
- Server virtualization
- Multi-core processors and blade servers
- Larger memories in servers
- Solid state disks
- Faster PCIe

Increased number and size of applications

The prolific nature of software applications drives the need for higher speed links. The big growth in all computing environments is due to the increasing number of software applications. The virtualization of information that was previously physical, along with an increased level of access to the data fundamentally drives bandwidth. From mobile access to email to cloud computing, data needs to be transported over longer distances at higher speeds. High-end applications such as databases and data warehousing continue to grow in size and are integrated into larger deployments. Corporations are integrating and sharing more applications that result in more links running at faster speeds.

Virtualized servers

In the past, each application would run on its own server and many servers were underutilized due to low demand from each application. With the adoption of server virtualization, many applications can run on a physical server and increase the bandwidth requirements of that server. For increased reliability and scalability, virtual servers are designed to migrate to different servers and require access to shared storage instead of direct-attached storage. Tens of applications per server drive higher bandwidth from each server and from each rack. This consolidation of applications and other technology builds a converged infrastructure and creates a need for faster links.

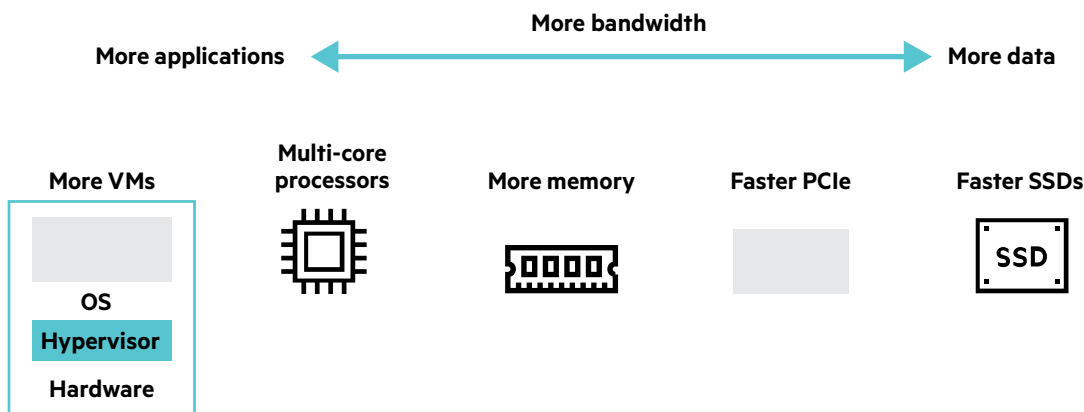


Figure 2. Faster links required

Multi-core processors

More processing power is needed per server to manipulate the large amounts of data that runs on multiple virtual servers and large applications. Enabling increased processing of the data is the relentless march of Moore's Law, which states that the number of transistors on an integrated circuit doubles every two years. One of the main uses of more transistors is in multi-core processors. These multi-core processors can drive bandwidth in excess of 50 Gbps. Some of the main multi-core processors include:

- Intel® Xeon® 7500 processor designed for 32-bit and 64-bit Microsoft® Windows® and Linux® applications. The 4-, 6-, and 8-core Nehalem EX uses Intel® Quick Path Interconnect (QPI) to drive up to 25 Gigabytes per second (Gbps) of bandwidth.
- Intel® Itanium® 9300 processor designed for 64-bit Microsoft Windows, Linux, and UNIX® applications. The 4-core Itanium 9300 also uses Intel QPI.
- AMD 8-core and 12-core Opteron 6000 series server processor with Hyper Transport 3.0 technology with 4.8 Giga transfers per second.
- IBM Power7 processors support 32 threads per processor and up to 1000 virtual machines per server. The 8-core Power7 processor uses RISC architecture for database and transaction-based workloads.
- SPARC T3 chip has 16 cores and supports up to 128 threads.

Larger memories in servers

Server memories grow quickly as long as terabytes (TBs) of data are supported on individual servers. With dual in-line memory modules (DIMMs), exceeding 10 GB per DIMM, most servers can support hundreds of GBs or even TBs of memory. As the capacity of these DIMMs continues to increase, memory will no longer be the limiting factor in deciding how many applications can be supported on a server.

The number of virtual machines that are supported on an individual server is usually limited by the memory capacity of the server or backup windows. Depending on how many GBs of memory an application needs and a server supports, a growing number of applications can be supported by a physical server. When memory limitations are overcome, a deployment may be limited by the backup server, which has a finite backup window. Depending on the deployment, the processing power, memory, or backup window restricts the number of virtual servers that can be aggregated in a physical server. Faster links will help curb this restriction, paving way for efficient servers and backup servers.

Solid state disks

In a high performing storage network, users must consider the incredible advances in SSDs. SSD capacities have grown to hundreds of GBs per drive and hundreds of thousands of I/Os per second (IOPS). SSDs decrease latency of storage from the millisecond range to the microsecond range to dramatically increase performance with low power consumption as well. SSDs can increase performance so that fewer servers are required to accomplish the same task.

Faster PCIe

Another advancement in computing architectures is based on evolution to the third generation of the PCIe 3.0 bus. PCIe 3.0 supports up to 64 Gbps of bandwidth and 4 ports of 16 Gbps FC. 16 Gbps FC HBAs enable higher utilization of the PCIe bus. Most servers are configured with 8 operational lanes while some are configured with 4 lanes. PCIe 3.0 switched to 128b/130b encoding to reduce the line rate. The line rate of PCIe 3.0 increases by only 3 Gbps while the data rate doubles from PCIe 2.0. 16 Gbps FC HBAs get the most use out of the PCIe bus and makes sure that the bandwidth bottleneck isn't in the HBA.

Table 3. PCIe data rates

PCIe	Number of lanes	Line rate (Gbps)	Data rate per lane (Mbps)	Directional bandwidth (Gbps)	Supported ports
PCIe 1.0	4	2.5	250	8	1 x 8 GFC port
PCIe 1.0	8	2.5	250	16	2 x 8 GFC or 1 x 16 GFC port
PCIe 2.0	4	5.0	500	16	2 x 8 GFC or 1 x 16 GFC port
PCIe 2.0	8	5.0	500	32	4 x 8 GFC or 2 x 16 GFC ports
PCIe 3.0	4	8.0	1000	32	4 x 8 GFC or 2 x 16 GFC ports
PCIe 3.0	8	8.0	1000	64	4 x 16 GFC port

Storage systems with 16 Gbps FC interfaces

Storage systems with 16 Gbps Fibre Channel interfaces are now available. Transitioning to 16 Gbps Fibre Channel infrastructure in data centers ensures that the fabric is no longer the bottleneck and that you get the most out of your Storage to support your consolidation and virtualization initiatives.

HPE Smart SAN—target-driven peer zoning for HPE 3PAR StoreServ Storage implementations

HPE Smart SAN is an innovative and standards-based approach through collaboration with HPE SAN partners that leverages Smart SAN enabled hosts, fabrics, and the HPE StoreServ 3PAR array to provide an end-to-end solution that makes SAN configuration and management simple and reduces the probability of errors through automated switch zoning. It is an application embedded in SAN components (switches, HBAs, and array) that enables the HPE 3PAR to orchestrate configuration, settings, and policies in a SAN, simply. HPE Smart SAN for HPE 3PAR through its target-driven peer zoning (TDPZ) feature enables customers to automate SAN configuration with peer zoning, resulting in the creation of fewer zones and enables configuration of zones in minutes and not in hours. Through automation, it reduces the probability of errors and potential downtime. HPE Smart SAN also supports standards based device registrations and diagnostics data collection for better configuration visibility and diagnostic purposes. For more information on HPE Smart SAN, you can access the free technical white paper [here](#).

Aggregation of traffic

The aggregate bandwidth from thousands of servers in large data centers can easily reach Terabits per second (Tbps) of bandwidth. Coming from mainframes, high-performance servers, midrange servers, blade servers, and volume servers—many streams of data add up to a torrent of storage traffic. When data backup and mirroring is added to the river of data, faster tools handle the job more efficiently. To better deliver Tbps of data, 16 Gbps FC links are more efficient than 8 Gbps FC or even 10GbE.

An example of a rack of servers with a Top-of-Rack (ToR) switch shows how traffic adds up quickly. ToR switches aggregate traffic from the 16 servers housed in one or more racks. If these servers average to 4 x 8 Gbps FC or 2 x 16 Gbps FC ports per server, then the equivalent of 32 x 16 Gbps FC ports are distributed to the two ToR switches. If a 4:1 oversubscription is adequate for application workloads, then either 8 x 16 Gbps FC or 16 x 8 Gbps FC ISLs are needed to deliver the necessary bandwidth to the core. The 16 Gbps FC ISLs is more cost-effective and easier to manage, when the additional cables, power consumption, and management of slower Fibre Channels are taken into consideration.

Think of blade servers as a mini-rack of servers with an embedded switch that acts as the aggregation point. Blade servers behave in a similar manner as a rack of servers; the main difference being that the rack has external links to the ToR switches instead of the backplane links in the blade server. The speed of external links in rack servers is often faster than the backplane links of blade servers. The backplane links of blade servers still aggregate large volumes of traffic. If a blade server has 16 server blades delivering 4 Gbps of backplane bandwidth, the blade server would need 4 x 16 Gbps FC or 8 x 8 Gbps FC ISLs to the core.

HPE recommends a core-edge topology for large SAN deployments. Edge connections from ToR switches or blade servers feed back to the core of the network where large modular switches route the data to other edge switches or end devices connected to the core. The large modular switches at the core of the network have many ISLs that aggregate traffic from the edge to other modular switches in the core. The core of the network will benefit from the use of 16 Gbps FC.

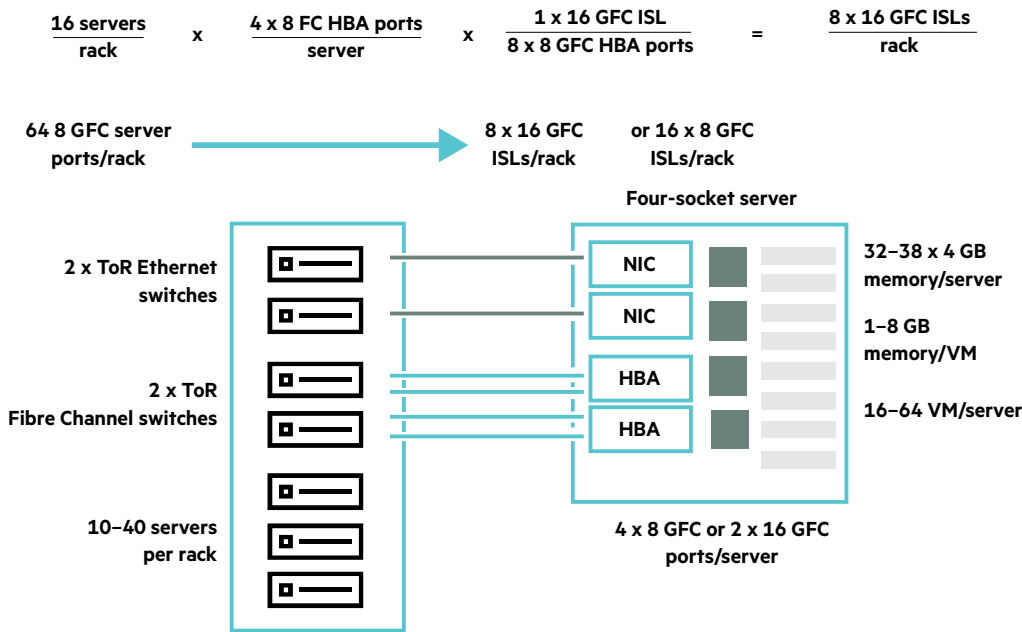


Figure 3. Traffic aggregation in a rack of servers

Applications of 16 Gbps FC

16 Gbps FC is designed for high-bandwidth applications and includes:

- ISLs
- High-bandwidth applications
- Data migration
- VDI
- SSD

High-bandwidth applications

The types of servers that will need the kind of performance that 16 Gbps FC can deliver run high-bandwidth applications whereas the majority of servers that use FC storage run database and enterprise-class applications. While database applications do not usually require large amounts of bandwidth when individual records are updated or read, servers need to be designed for demanding tasks such as backup and data mining, during which every record may be copied or queried. Backup and recovery applications are the high-water marks for which servers need to be designed.

Another class of applications that benefit from 16 Gbps FC is streaming I/O. A single I/O from these applications can transfer a block of data that is several orders of magnitude larger than blocks in general purpose file systems. A single I/O can take minutes or hours to complete while controllers and drives are sending out sequential reads or writes as fast as they can.

Data migration

Another application of 16 Gbps FC links in the data center is between data centers, storage arrays, or clouds. During data center consolidation, disaster recovery, and equipment changes, users often have a need to migrate Terabytes or even Petabytes of data between storage arrays. The time to transfer large blocks of data is often limited by the speed of the links connecting the devices instead of processors or controllers that may limit the throughput during normal processing. To enable migration of applications to different locations or clouds, the data needs to be present in both locations before the application can migrate, so 16 Gbps FC can facilitate these data migrations and mirroring.

Virtual Desktop Infrastructure

VDI is a growing trend in environments where virtual desktops in the data center are sent to users on a variety of devices. VDI has the advantage of centralized management, allowing applications and hardware to be upgraded easily in the data center and virtually shipped around the world. VDI has large bandwidth requirements when large numbers of users log into their virtual desktops at the same time. This spike in activity leads to long startup times unless high-performance VDI systems are used. 16 Gbps FC is one of the many components that can help improve performance at these critical initialization times.

Solid state disks

As mentioned earlier, SSDs are enabling a new level of performance in latency-sensitive applications. With lower latency and higher IOPs than traditional storage arrays, 16 Gbps FC interfaces to SSDs are expected to improve the bandwidth density of their front panels by doubling the throughput of their ports. SSDs have been applied to many high-bandwidth situations such as online gaming, in which bandwidth requirements of 50 Gbps have already been reached. With the price of SSDs dropping quickly, SSDs should be able to address many more applications in which performance is more important than capacity.

16 Gbps FC vs. FCoE

There are mixed opinions, with some experts stating that Ethernet is going to take over the networking world and that FCoE is going to “kill” FC. The reality is that both technologies will coexist for many years to come and have appropriate applications. For servers that aren’t pushing over 10 Gbps of LAN and SAN traffic, 8 Gbps FC HBAs or 10GbE Converged Network Adapters (CNAs) might be the most appropriate choices. For environments where combined Ethernet and FC traffic exceed 10 Gbps, an 8 Gbps FC HBA or 16 Gbps FC HBA and a 10GbE NIC might be more efficient than two 10GbE CNAs.

Considering that a 16 Gbps FC HBA delivers 40 percent more throughput than a 10GbE CNA, fewer ports may be consumed on blade servers or high-end servers that require tens of Gbps of throughput.

Demartek key findings

Transaction-intensive database workloads are sensitive to storage I/O latency and, in the case of particularly busy applications, storage bandwidth. This evaluation’s online transaction processing (OLTP) workload, deployed on an HPE 3PAR StoreServ 8450 All-Flash Storage System, saturated a single port of 8 GbFC bandwidth. This I/O limit drove up latency due to the queuing of I/O requests and placed a hard limit on the number of database transactions the server could process per second. By replacing 8 GbFC components with GEN 5 16 GbFC, we doubled the transaction rate by improving bandwidth and IOPs more than 70 percent while cutting latency by a similar percentage. [Click](#) for full report.

We recommend

It’s easy to understand that a link that is twice as fast as the much slower link can get more work done. While many data environments are not utilizing all the capabilities of 16 Gbps FC links just yet, over the next few years, traffic and applications will grow to fill the capacity of 16 Gbps FC. The refresh cycle for networks is often longer than that of servers and storage, so 16 Gbps FC will remain in the network for years. With the addition of more virtual machines to a physical server, performance levels can quickly escalate beyond the levels supported by 8 Gbps FC. To future-proof deployments, 16 Gbps FC should be considered to be the most efficient way to transfer large amounts of data in data centers. With HPE B-series trunking technology at 16 Gbps FC, users can get up to 128 Gbps of FC performance, which delivers more bandwidth for the power and cost of the deployment.

16 Gbps FC is slated to be the best performer in several environments. It can reduce the number of ISLs in the data center, or migrate a large amount of data for array migration or disaster recovery. High-performance applications, such as VDI, which use SSDs or require high bandwidth, are ideal for 16 Gbps FC. 16 Gbps FC uses more efficient 64b/66b encoding to increase efficiency and transmitter training to tune links for best performance. 16 Gbps FC combines the latest technologies in an energy-efficient manner to enable the highest performing SANs in the world.

Top four reasons to partner with HPE

We have in-house capabilities to address changing data center needs. We understand the demands on your business to speed up delivery while lowering costs and maintaining efficiency.

1. Valuable experience

We have the experience of simplified SAN by combining FC performance with a complete SAN management system, allowing for the fact that better technology only means speedier and improved performance with 16 Gbps FC.

2. Innovation at every step

A converged infrastructure environment provides you with a blueprint for a data center that can handle the growing volume of data that is exchanged. We build innovation with the future in mind. Whether your needs involve email/database scalability, server/storage consolidation, or automated backup/disaster recovery—our 16 Gbps FC solutions can help you accomplish your tasks while providing you an accelerated ROI and enabling a virtualized setup that achieves better business results.

3. Seamless integration

Whether it is your storage array migration, disaster recovery, VDI, and ISL applications that demand higher speed to transmit large chunks of data, our 16 Gbps FC solutions can help seamlessly integrate new technology with the existing applications, ensuring your data needs are never compromised.

4. Proven success

As one of the worldwide leaders in SAN products and solutions, our storage solutions can help you enable an adaptive infrastructure that is flexible, giving your data environment the scalability and performance speeds that are required for critical operations.

Learn more at

hpe.com/storage/san



Sign up for updates

© Copyright 2011, 2013, 2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

AMD is a trademark of Advanced Micro Devices, Inc. Intel, Intel Itanium, and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. UNIX is a registered trademark of The Open Group. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

4AA3-6611ENW, August 2016, Rev. 2