DATA CENTER

The Benefits and Application of 16 Gbps Fibre Channel

Discusses 16 Gigabits per second (Gbps) Fibre Channel (FC) and how it improves throughput in Storage Area Networks to reduce Inter-Switch Link counts, improve application performance, ease management, and reduce power consumption per bit.

BROCADE

Contents

INTRODUCTION

The latest speed developed by the T11 technical committee that defines Fibre Channel interfaces is 16 Gbps Fibre Channel. The Fibre Channel industry is doubling the data throughput of 8 Gbps links from 800 Megabytes per second (MBps) to 1,600 MBps with 16 Gbps FC. 16 Gbps FC is the latest evolutionary step in Storage Area Networks (SANs) where large amounts of data are exchanged or high performance is a necessity. From Host Bus Adapters (HBA) to switches, 16 Gbps FC will enable higher performance with lower power consumption per bit—the performance required by today's leading applications.

The benefits of a faster technology are easy to see. Data transfers are faster, fewer links are needed to accomplish the same task, fewer devices need to be managed, and less power is consumed when 16 Gbps FC is used instead of 8 Gbps or 4 Gbps. 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. 16 Gbps FC is keeping pace with other technology advances in the data center.

When high bandwidth is needed, 16 Gbps FC should be deployed. Applications in which bandwidth demands are high include storage array migration, disaster recovery, Virtual Desktop Infrastructure (VDI) and Inter-Switch Links (ISLs). The first place that new speeds are usually needed in SANs is in ISLs in the core of the network and between data centers. When large blocks of data need to be transferred between arrays or sites, a faster link can accomplish the same job in less time. 16 Gbps FC is designed to assist users in transferring large amounts of data and decreasing the number of links in the data center.

NOTE: The abbreviation "Gbps Fibre Channel" is sometimes abbreviated even further in this paper to "GFC" for convenience.

OVERVIEW OF 16 GBPS FC

Offering considerable improvements from the previous Fibre Channel speeds, 16 Gbps FC uses 64b/66b encoding, retimers in modules, and transmitter training as outlined in Table 1. 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, Electronic Dispersion Compensation (EDC) and transmitter training to improve backplane links. The combination of these technologies enables 16 Gbps FC to provide some of the highest throughput density in the industry.

Table 1. Fibre Channel 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

While 16 Gbps FC doubles the throughput of 8 Gbps FC to 1600 MBps, the line rate of the signals only increases to 14.025 Gbps because of a more efficient encoding scheme. Like 10 Gbps FC and 10 Gigabit Ethernet (GbE), 16 Gbps FC uses 64b/66b encoding, that is 97% efficient, compared to 8b/10b encoding, that is only 80% efficient. If 8b/10b encoding was used for 16 Gbps FC, the line rate would have been 17 Gbps and the quality of links would be a significant challenge because of higher distortion and attenuation at higher speeds. By using 64b/66b encoding, almost 3 Gbps of bandwidth was dropped from the line rate so that the links could run over 100 meters of Optical Multimode 3 (OM3) fiber. *By using 64b/66b encoding, 16 Gbps FC improves the performance of the link with minimal increase in cost.*

To remain backward compatible with previous Fibre Channel speeds, the Fibre Channel Application-Specific Integrated Circuit (ASIC) must support both 8b/10b encoders and 64b/66b encoders. As seen in Figure 1, a Fibre Channel ASIC that is connected to an SFP+ module has a coupler that connects to each encoder. The speed-dependent switch directs the data stream toward the appropriate encoder depending on the selected speed. During speed negotiation, the two ends of the link determine the highest supported speed that both ports support.

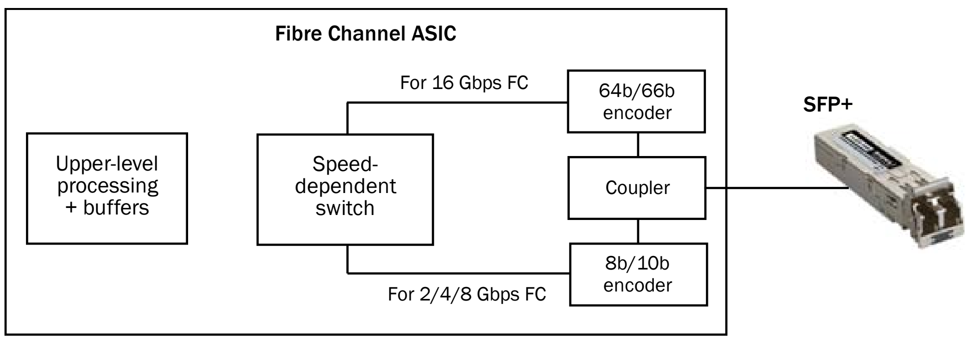

Figure 1. Dual codecs

Figure 2 shows how a port that supports 16 Gbps FC interoperates with 2/4/8 Gbps FC ports and 16 Gbps FC ports. In the top half of the illustration, a port that supports 2/4/8/16 Gbps FC is connected to a port that supports only 2/4/8 Gbps FC. Since the fastest speed that both ports support is 8 Gbps FC, the port runs at 8 Gbps FC through the 8b/10b encoder. If a link connects two 16 Gbps FC-capable ports, as shown in the bottom half of Figure 2, the encoders will use the 64b/66b encoders and run at 16 Gbps FC. The 16 Gbps FC ASIC is designed to maximize investments in Fibre Channel technologies by being backward compatible with existing infrastructure.

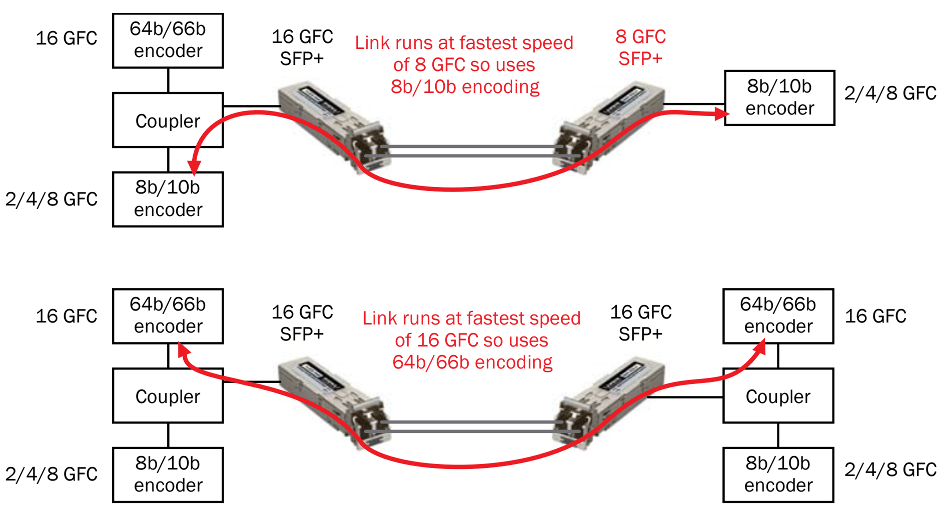

Figure 2. Negotiated speeds

The second technique that 16 Gbps FC uses to improve link performance is the use of retimers or Clock and Data Recovery (CDR) circuitry in the SFP+ modules. The most significant challenge of standardizing a high-speed serial link is developing a link budget that manages the jitter of a link. Jitter is the variation in the bit width of a signal due to various factors, and retimers eliminate most of the jitter in a link. By placing a retimer in the optical modules, link characteristics are improved so that the links can be extended for optical fiber distances of 100 meters on OM3 fiber. The cost and size of retimers has decreased significantly so that they can now be integrated into the modules for minimal cost.

The 16 Gbps FC multimode links were designed to meet the distance requirements of the majority of data centers. Table 2 shows the supported link distances at multiple speeds over multimode and single-mode fiber—16 Gbps FC was optimized for OM3 fiber and supports 100 meters. With the standardization of OM4 fiber, Fibre Channel has standardized the supported link distances over OM4 fiber and 16 Gbps FC can 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. This wide range of supported link distances enables 16 Gbps FC to work in a wide range of environments.

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

Speed	OM1 Link Distance 62.5 um core 200 MHZ*km	OM2 Link Distance 50 um core 500 MHz*km	OM3 Link Distance 50um core 2000 MHz*km	OM4 Link Distance 50um core 4700 MHz*km	OS1 Link Distance 9um 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

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

Another important feature of 16 Gbps FC is that it uses transmitter training for backplane links. Transmitter training is an interactive process between the electrical transmitter and receiver that tunes lanes for optimal performance. 16 Gbps FC references the IEEE standards for 10GBASE-KR, which is known as "Backplane Ethernet," for the fundamental technology to increase lane performance. The main difference between the two standards is that 16 Gbps FC backplanes run 40% faster than 10GBASE-KR backplanes for increased performance.

The Benefits of Higher Speed

The benefits of faster tools are always the same—more work in less time. By doubling the speed, 16 Gbps FC reduces the time to transfer data between two ports. When more work can be done by a server or storage device, fewer servers, HBAs, links, and switches are needed to accomplish the same task. The benefits of 16 Gbps FC add up and include:

- Reduced number of links, HBAs and switch ports to do the same workload
- Reduced power consumption per bit
- Easier cable management

Reduced Number of Links

As with other Fibre Channel speeds, the first application of new Fibre Channel speeds is on ISLs between switches. Large fabrics are composed of many switches that are connected via multiple ISLs. Reduction of the number of ISLs between switches is a key benefit of each higher speed. Brocade switches will continue to 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 table in Figure 3 shows a simple comparison of the number of links in an 8 Gbps FC fabric and a 16 Gbps FC fabric. The higher-speed links of 16 Gbps FC eliminate tens or hundreds of ports from a comparable 8 Gbps FC fabric. The real savings occur when the number of HBAs, switches, and end devices can be decreased with the higher performance of 16 Gbps FC. In the example in Figure 3, a Top of Rack (ToR) switch needs 100 Gbps of bandwidth so the user needs 8 x 16 Gbps FC ISLs instead of 16 x 8 Gbps FC ISLs. Similar comparisons between 16 Gbps FC ISLs and 8GFC ISLs are given in the table in Figure 3 to show how fewer ports and links are needed at 16 Gbps FC.

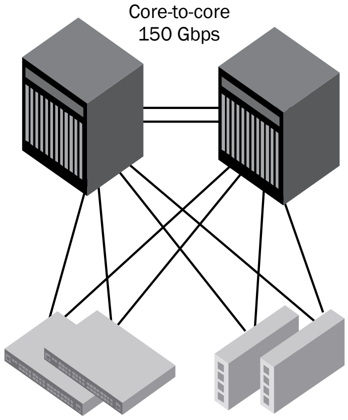

	8 Gbps FC Links	16 Gbps FC Links
ISLs from ToR switch to core	16	8
ISL from blade switch to core	10	5
Core to core	24	12
Total ISLs	**50**	**25**
Total ports	**100**	**50**

Figure 3. Network design implications

Reduced Power Consumption per Bit

Besides the reduction in equipment that cuts power consumption dramatically, 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 Total Cost of Ownership (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 16 Gbps FC SFP+s consuming 0.75 watts of power while 8GFC SFP+ consuming 0.5 watts of power. These estimates show that a 16 Gbps FC link consumes 25% less power than 2 x 8 Gbps FC ports.

Easier Cable Management

If fewer links are needed, cable management becomes simpler. Managing cables behind a desktop or home entertainment center is difficult enough, but managing hundreds of cables from a single switch or bundles of cable from a server can be horrendous. The reduction of cables means less troubleshooting and recabling. The cost of cabling is significant and users can pay over $300/port in structured cabling environments. Reducing the number of links by using fast 16 Gbps FC links reduces the work and cost involved in cable management.

Brocade has teamed with Corning Cabling Systems to provide incredibly dense solutions for cable management. To reduce the bulk of cables, Corning offers uniboot cables, which combine 2 fibers into one cord and then 12 fibers into one ribbon cable, as shown in Figure 4. The LC to MTP cable harnesses reduce cable bulk and utilize compact fiber ribbons. Corning also provides very dense patch panels for MTP and LC connectors. Brocade and Corning offer excellent solutions for all types of cable management.

Figure 4. Uniboot LC to MTP cable harness

Summary of Benefits

The end result of 16 Gbps FC is that there are fewer links, fewer cables, fewer ports and less power for the same performance. Figure 5 compares one 16 Gbps FC link to two 8 Gbps FC links. The most significant benefit of 16 Gbps FC ports will be the need for fewer HBAs and switch ports connected to these media.

Figure 5. Comparison of media

The Need for Speed

The principle benefit of 16 Gbps FC is the increased performance at double the data rate of 8 Gbps FC and 40% faster than 10 GbE. A question that often arises when speeds are doubled is this: is this new bandwidth really needed and where?

Multiple technology advances are driving the need for 16 Gbps FC including:

- Increased number and size of applications
- Server virtualization
- Multi-core processors and blade servers
- Larger memories in servers
- Solid State Disks (SSDs)
- Faster PCIe

More Applications

The prolific nature of software applications drives the need for higher speed links. The big growth in all computing environments is the increasing number of software applications. The digitization of information, that was previously physical, and increased access to the data are what fundamentally drives bandwidth. From mobile access to e-mail 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.

Server Virtualization

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 (DAS). Tens of applications per server drive higher bandwidth from each server and from each rack. This consolidation of applications and other technology advances create a need for faster links, as shown in Figure 6.

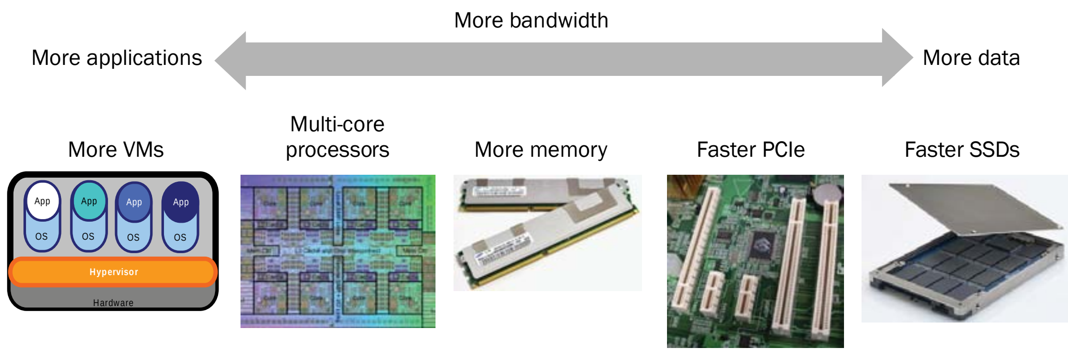

Figure 6. Faster links

Multi-core Processors

More processing power is needed per server to manipulate the large amounts of data that run 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[1].
- Intel's Itanium 9300 processor designed for 64-bit Microsoft Windows, Linux, and Unix applications. The 4-core Itanium 9300 also uses Intel QPI[2].
- AMD 8-core and 12-core Opteron 6000 series server processor with HyperTransport 3.0 technology with 4.8 GigaTransfers per second.
- IBM Power7 processors support 32 threads per processor and up to 1000 virtual machines per server. The 8-core Power7 processor uses a RISC architecture for database and transaction-based workloads[3].
- Sun UltraSparcT3 chip has 16 cores and supports up to 128 threads.

More Memory

The memory in servers is growing quickly until terabytes (TBs) of data are supported on individual servers. With Dual Inline Memory Modules (DIMMs) exceeding 10 GB per DIMM, most servers can support hundreds of Gigabytes (GBs) or even TBs of memory. As the capacity of these DIMMs continues to increase, memory will no longer be the limiting factor in 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 Gigabytes 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 limited backup windows. Depending on the deployment, the processing power, memory, or backup window limits the number of virtual servers that can be aggregated in a physical server.

Solid State Drives

If high performance in storage is required, users must consider the incredible advances in solid state drives (SSDs). SSD capacities have grown to hundreds of GBs per drive and hundreds of thousands of I/Os per second (IOPS). The RamSan-630 is the latest SSD-based storage system from Texas Memory Systems and supports 500,000 IOPS and 8 GBps (64 Gbps) per array. 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 advance in computing architectures is based on the third generation of the Peripheral Component Interconnect Express (PCIe) 3.0 bus. PCIe 3.0 can support up to 64 Gbps of bandwidth and 4 ports of 16 Gbps FC, as seen in Table 3. 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 the bandwidth bottleneck isn't in the HBA.

Table 3. PCIe data rates

	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 - 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

When all of these technology advances are considered, the need for 16 Gbps FC becomes more apparent. An example of a high-performance computer shows how these technologies are combined to form incredibly fast compute systems that support hundreds of Gbps of I/O. The HP Superdome is one of many computers whose architecture can support multiple 16 Gbps FC HBAs. As seen in Figure 7, the HP Superdome can support tens of 16 Gbps FC HBAs when tens of processors are in use. These high-performance computers are driving a new level of application performance.

	32 CPU	**64 CPU**	**128 CPU**
Memory/partition	512 GB	1 TB	2 TB
Physical partitions	4	8	16
Memory bandwidth	64 GBps	128 GBps	256 GBps
2.133 MBps PCI slots	24	48	96
Peak I/O	368 Gbps	736 Gbps	1472 Gbps
Number of 16 Gbps FC HBAs*	11	22	44

* Based on 150, 300, and 600 Gbps of bandwidth to the SAN

Figure 7. HP Superdome 128 CPU capabilities

Aggregation of Traffic

The aggregate bandwidth from hundreds or 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 10 GbE.

An example of a rack of servers with a Top of Rack (ToR) switch helps show how traffic adds up quickly. As seen in Figure 8, ToR switches aggregate traffic from the 16 servers housed in one or more racks. If these servers average 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 deemed 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 will probably be more cost effective and easier to manage when the additional cables, power consumption, and management are taken into consideration.

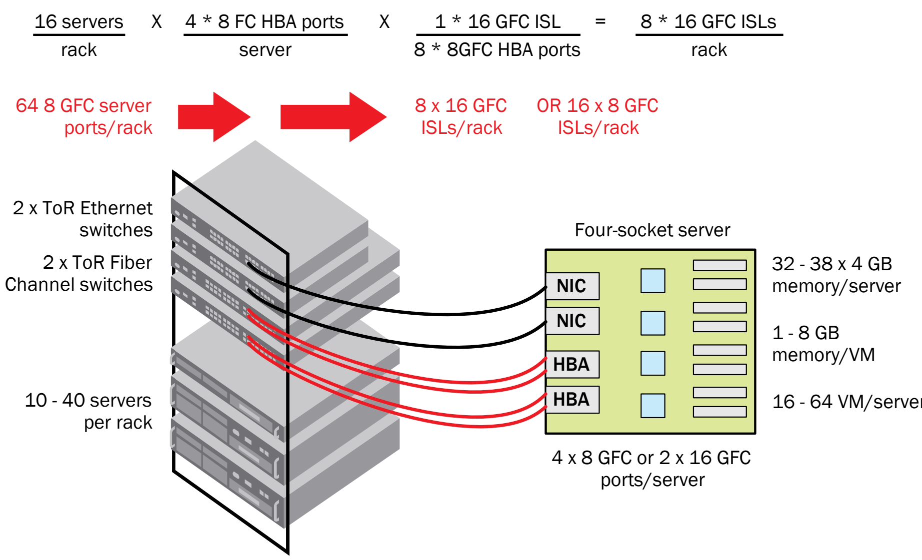

Figure 8. Traffic aggregation in a rack of servers

NOTE: Special mention of the bandwidth requirements for mainframe computers is worthwhile. High Performance FICON for System Z (zHPF) can drive 8 Gbps FC links at full capacity[4] and will eventually drive 16 Gbps FC links to full capacity. Mainframes have been improved and tuned for decades and these will continue to be major workhorses for the data center.

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 switches; the main difference is 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. Blade server backplanes usually run at 3.125 or 6.25 Gbps with some 10GBASE-KR (10 GbE backplane) backplane links being developed now. 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 needs 4 x 16 Gbps FC or 8 x 8 Gbps FC ISLs to the core.

Brocade 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.

Applications of 16 Gbps FC

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

- ISLs (discussed in the previous section)
- High-bandwidth applications
- Data migration
- Virtual Desktop Infrastructure (VDI)
- Solid State Disks (SSD)

High-Bandwidth Applications

The types of servers that need 16 Gbps FC HBAs run high-bandwidth applications. The majority of servers that use Fibre Channel 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. Table 4 shows the time required to transfer large amounts of data at 1,600 MBps with 16 Gbps FC. When time is money, 16 Gbps FC is better.

Table 4. Data migration examples

Data Size	Time to Transfer Data at 1600 MBps
100 GB	1 minute
1 TB	10 minutes
10 TB	1 hour, 45 minutes
100 TB	17 hours
1 PB	1 week

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

Some pundits express the opinion that Ethernet is going to take over the networking world and that Fibre Channel over Ethernet is going to “kill” Fibre Channel. 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 Local Area Network (LAN) and SAN traffic, then 8 Gbps FC HBAs or 10 GbE Converged Network Adapters (CNAs) might be the most appropriate choices. For environments where combined Ethernet and Fibre Channel traffic exceed 10 Gbps, an 8 Gbps FC HBA or 16 Gbps FC HBA and a 10 GbE NIC might be more efficient than two 10 GbE CNAs. Considering that a 16 Gbps FC HBA delivers 40% more throughput than a 10 GbE CNA, fewer ports may be consumed on blade servers or high-end servers that require tens of Gbps of throughput.

Summary

Speed wins! It’s easy to understand that a link that is twice as fast as a slower link can do more work. While many environments won’t use the full extent of a 16 Gbps FC link 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 more virtual machines being added 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 Brocade trunking technology at 16 Gbps FC, users can get up to 128 Gbps of Fibre Channel performance, which delivers more bandwidth per power and cost.

16 Gbps FC will 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. In conclusion, 16 Gbps FC combines the latest technologies in an energy-efficient manner to enable the highest performing SANs in the world.

References

1. Intel QuickPath Architecture, http://www.intel.com/technology/quickpath/whitepaper.pdf
2. Intel Itanium 9300 Processor Raises Bar for Scalable, Resilient Mission Critical Computing, http://www.intel.com/pressroom/archive/releases/20100208comp.htm
3. IBM Unveils New POWER7 Systems to Manage Increasingly Data-Intensive Services, http://www-03.ibm.com/press/us/en/pressrelease/29315.wss
4. IBM System z10 FICON Express8 Channel Performance, ftp://public.dhe.ibm.com/common/ssi/sa/wh/n/zsw03127usen/ZSW03127USEN.pdf