FIBRE CHANNEL SOLUTION GUIDE

FIBRE CHANNEL INDUSTRY ASSOCIATION

State of the Fibre Channel Industry

Today's data explosion presents unprecedented challenges incorporating a wide range of application requirements such as database, transaction processing, data warehousing, imaging, integrated audio/video, real-time computing, and collaborative projects. For nearly a decade storage area networks (SANs) have become mainstays for companies looking to increase storage utilization and manageability while reducing costs. SANs represent a topology for connecting storage assets directly to the network and establishing a peer-to-peer server/storage implementation and solve multiple issues for enterprises with data centers to remote offices.

As the volume and criticality of data grow, companies need efficient, scalable solutions for making data available to servers, applications, and users across the enterprise. By providing a network of storage resources to servers, Fibre Channel SANs uncouple storage from individual platforms, allowing data transfer among all nodes on the storage network.

Fibre Channel is an ideal solution for IT professionals who need reliable, cost-effective information storage and delivery at fast speeds. With development starting in 1988 and ANSI standard approval in 1994, Fibre Channel is a mature, safe solution for 1Gb, 2Gb, 4Gb, 8Gb and 16Gb communications, providing an ideal solution for fast, reliable mission-critical information storage and retrieval for today's data centers. Another new Fibre Channel interface development is Fibre Channel over Ethernet (FCoE), which leverages the millions of man-hours of Fibre Channel upper layer management and driver development and its extensive installed base, now shipping at 10GFCoE and coming soon industry-wide at 40GFCoE and 100GFCoE.

THE SPEED YOU NEED

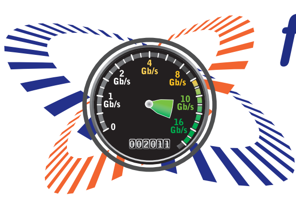

for Big Data and Flexibility

16GFC & FCoE

Articles Included:

- The FCIA Maps FC to FCoE to Bandwidths Needed Now and Later!
- Twice as Fast Is More Than Two Times Better
- Unified Fabric and FCoE
- Unified Fabric Management

THE FCIA MAPS FC TO FCoE TO BANDWIDTHS NEEDED NOW AND LATER!

Skip Jones - Chairman, FCIA (Fibre Channel Industry Association)

The heart and soul of any technology, and the industry association that stewards the technology, is its technology roadmap. Just like the term suggests, a roadmap shows the history and future of a technology. It is a guide to where it has been, where it is going and when it is going to get there. The three primary audiences for a technology roadmap are the user base that deploys the technology, the development, manufacturing and distribution base that supplies the technology, and the industry standards bodies that develop standards for the technology.

A consistently trustworthy roadmap provides the user with a planning document. Additionally, the roadmap provides the user with confidence that their investments in the technology will be preserved into the foreseeable future. The roadmap shows that the technology has legs to run with and thereby ensures their investments today are future-proofed for tomorrow.

An accurate roadmap provides a reliable guide for suppliers to plan their product development and release cycles based upon the features and timing of the technology migration reflected in the roadmap.

A dependable and responsible roadmap provides standards bodies a planning cookbook by which they can initiate and complete standards within the timeframe defined by the roadmap. The roadmap also directs suppliers on when to begin product development using said technology. The supplier's development efforts are based upon open standards that are technically complete. Some technology developments are required building blocks for further product development. For example, lasers in optical modules need to be developed before the modules can be developed that will eventually be used in a switch or host bus adapter. With a solid roadmap and standards, multiple companies can develop products in parallel that will eventually interoperate when they reach the market.

So how does a technology roadmap become a responsible, reliable, trustworthy and consistently accurate planning document? The short answer is that it takes time and commitment, and perhaps most of all a relentless effort to match its technology migration to the requirements of the applications that roadmapped technology will be serving at any point in time. It takes years for the roadmap to have a sufficiently deep history that has year-in and year-out kept its promise to become credible. It must be a stable and consistent document that does not frequently change and reset expectations in the industry. And it must provide the speeds and feeds needed by the applications it will be used in.

A changing roadmap causes confusion and could cause faulty planning from user and supplier based upon an erroneous, ever-changing inaccurate roadmap. In order to avoid loss of credibility and trust from standards creators, technology suppliers and end users, it simply must have a rich history of being solidly accurate in its past forecasts.

One of the best industry examples of a roadmap that meets this proven reliable, trustworthy criterion is the FCIA roadmap. Since 1997, the FCIA roadmap has been spot-on with its mapping of Fibre Channel speeds. In addition to the Fibre Channel speeds, the FCIA has also mapped the timeline and speed migration for FCoE. FCIA success in delivering 14 years of accurate roadmaps come from the seriousness FCIA takes in this huge responsibility and obligation to the industry.

Continued...FCIA Maps FC to FCoE

A key attribute to this success is FCIA's unwavering commitment to understanding the applications that FC technology (and now FCoE as well) will be relied upon to serve, and retain the flexibility to change with the times the market needs served. For instance, in the late '90s FC was almost entirely used as a device-level interface for high-end disk drives, and somewhat for high-end tape products. The interface technology was expected to migrate at the same rhythm as disk drives migrated their capacity, which directly increases their data density which directly enables and demands higher transfer rates.

As the new millennium approached the technology was being used to connect the new advent of "boxes" of these disk and tape drives, which created the need for higher and higher bandwidths as well as the notion of "core" bandwidths coming out of the storage box that concentrated the additive speed transfer rates of the devices inside the box.

Then as the early part of the millennium gave birth to geographically distributed enterprise data centers that needed to access and share massive amounts of data, these storage boxes started being place onto a switched network called SANs (Storage Area Networks), thus pushing the FC technology into more of a switched network technology and away from its previous focus as a direct-connect device-level I/O technology using FC-loop as its basis. Also, other device-level technologies were making their entrance as device I/O solutions such as SATA (serial ATA) and SAS (Serial Attached SCSI), allowing the FC community to focus more of its efforts on the blossoming SAN market.

Then a few years ago enterprise data centers began re-centralizing and becoming less and less distributed geographically, creating a need for higher shared I/O bandwidths and new topologies of storage networks. One of Fibre Channel's key attractions throughout all this history of changes and migration has been its solid backward compatibility. At the software driver and management layers FC looks much the same to the software interface that interconnects with it the same now as it did in 2001, allowing all of these intensely complicated software interfaces to migrate unscathed and remain portable. It is these "untouchable" upper software management and driver layers that needed to be ported into the new decentralization of data centers and the resulting consolidation of storage and servers.

From this advent came the notion of porting all of these storage FC-centric driver and management software layers untouched onto Ethernet cabling at the physical layer, and FCoE (Fibre Channel over Ethernet) was born. FCoE looks just like FC to these entrenched upper layers but simply operates over Ethernet physical infrastructure. Also making FCoE extremely attractive is its ability to be used in conjunction with traditional client Ethernet network traffic all on the same Ethernet physical infrastructure. Using the same Ethernet physical infrastructure for both traditional client network traffic and FC storage traffic via FCoE is commonly referred to as a "Converged Network".

Also the new increased consolidation of storage increased the need for much higher Fibre Channel physical layer bandwidth since more and more access was being demanded of this increasingly shared storage-only interconnect, so it is no wonder that 8GFC (8 gigabit per second Fibre Channel) has become the new FC standard shipping in FC components, but also operate at backward compatible 4GFC and 2GFC speeds.

Now, looking ahead, we see the need for speed and IOPS going through the roof with the advent of new data center approaches such as Cloud Storage and Cloud Computing. Whether private or public, these "clouds" serve as more and more of a "shared" bandwidth which was being highly accessed even before it was being shared, thus the new expectation of it being shared places even more demand on the storage network. It is more than coincidence that

Continued...FCIA Maps FC to FCoE

a SAN is schematically illustrated with the image of a "cloud", as most networks are shown. The entire premise of SAN stems from the same mindset from which cloud storage stems. Many servers sharing the same storage (SAN), or many users sharing the same servers and storage (Cloud), both share many similarities.

To meet the increased need for storage access via shared networks required by clouds, as the FCIA roadmap shows, 16GFC is coming this year for storage-only SANs and 40GFCoE will be more and more used in converged networks, especially for Inter-Switch Links (ISLs) which usually require higher bandwidth than the numerous "edge" connections all feeding into the ISL, followed by 32GFC edge and ISL and 100GFCoE ISL.

Another industry move that demands higher than ever bandwidth is the increased usage of Solid State Disk (SSD). Whilst the usually less important WRITE times of these devices are considerably slower than rotational/actuator accessed media disk drives, their READ times are phenomenally faster, and most databases are needed to be read from faster than written to, so banks of SSDs are becoming more prevalent in high-end data centers. Regardless of what their device-level I/O is, they are interfaced as boxes of SSDs using the tried and true SAN interface of Fibre Channel or FCoE because upper-layer driver and management software "looks" at this storage system the same as it does good old proven tried and true Fibre Channel storage.
FCIA has a Roadmap Committee that is closely associated with INCITS T11.2 Task Group, the standards body that defines Fibre Channel speeds. Since FCIA meets at the T11 meetings, and its roadmap committee include many of the key T11.2 standards engineers as well as key Fibre Channel supplier corporate and technical marketing experts, the resulting roadmap is the refined product of an intense iterative process that pinpoints highly attractive market propositions balanced with sound engineering feasibility. The end result is an official FCIA roadmap and set of MRDs (Marketing Requirement Documents) that becomes T11.2's map of speeds and timelines. The MRDs define sets of features and benefits that are not only feasibly doable within the roadmap timelines, but it also results in actual products delivered in the prescribed timeframe that realize massive market success.

T11.2, like any standards body, is allergic to wasting time developing standards that never see the light of day in successful markets. That is one key reason that FCIA's roadmap, different from other industry roadmaps, takes great pains in accurately defining when a technically stable standards document is required to enable a specific speed migration and products based upon that speed.

FCIA's defined process of roadmap development has over the years earned the trust from T11.2 to the point that its MRDs and resulting roadmap become INCITS documents embedded in the standards development process. The roadmap ensures that what goes down on paper for official standards are within the MRD's guidelines and encompassed true market requirements.

This successful FCIA/T11 process of roadmap development and relentless execution results in reliable, relevant standards. The resulting standards are stable and ready in time for suppliers to begin their development. They are standards that meet feature/benefit criteria and guarantee functionality, cost, compatibility, power, length, and other components for a successful market. The user benefits by having a wide selection of products based upon open standards in a timeframe that meets the user's demands.

FCIA Roadmap Committee members as well as many T11.2 developers are also closely associated with Ethernet standards bodies such as IEEE 802. This close association ensures that our roadmap for FCoE is consistent with Ethernet migration, and visa versa.

Continued...FCIA Maps FC to FCoE

FCIA's Roadmap, version V13 (attached), is the latest descendent of a long successful history of the FCIA roadmap and can be found at: http://www.fibrechannel.org/roadmaps. It maps the doubling of Fibre Channel speeds from 1GFC (Gigabits per second Fibre Channel), 2GFC, 4GFC all the way out to 512GFC in factors of 2 GFC for edge connectivity. Each doubling of speed has taken about 3 years to complete and the 32GFC standard is expected to be stable in 2012. It also maps FC and FCoE ISL's (Inter-Switch Links) out to 1TFC (1 Terabit/s Fibre Channel) and 1TFCoE (1Terabit/s Fibre Channel over Ethernet). The V13 Roadmap also pinpoints standard stability and general market availability for 16GFC and 32GFC edge connectivity (16GFC in 2011 and 32GFC in 2014). The third page of the FCIA Roadmap, v13, is the FCoE-only migration of bandwidth. This roadmap shows the long legs that Fibre Channel has going into the future.

Other important elements defined in the roadmap include backward compatibility. For instance, just like 1GFC, 2GFC, 4GFC, and 8GFC edge connectivity, 16GFC and 32GFC are required to be backward compatible at least two generation. These speeds are auto-negotiated with no user intervention required, - i.e., 16GFC will automatically run at 4GFC and 8GFC, whilst 32GFC will automatically run at 8GFC and 16GFC. This important level of backward compatibility has been and will continue to be a major benefit in Fibre Channels continued success.

Conclusion

Technology roadmaps are important for guiding users, suppliers and standards bodies to a technology destination in a coordinated fashion. The FCIA Roadmap's consistency, reliability and accuracy has given these audiences a plan they can rely on to show them directions to success.

Fibre Channel Roadmap

Product Naming	Throughput (MBps)	Line Rate (GBaud)†	T11 Spec Technically Completed (Year) ‡	Market Availability (Year) ‡
1GFC	200	1.0625	1996	1997
2GFC	400	2.125	2000	2001
4GFC	800	4.25	2003	2005
8GFC	1600	8.5	2006	2008
16GFC	3200	14.025	2009	2011
32GFC	6400	28.05	2012	2014
64GFC	12800	TBD	2015	MARKET DEMAND
128GFC	25600	TBD	2018	MARKET DEMAND
256GFC	51200	TBD	2021	MARKET DEMAND
512GFC	102400	TBD	2024	MARKET DEMAND

- *"FC" used throughout all applications for Fibre Channel infrastructure and devices, including edge and ISL interconnects. Each speed maintains backward compatibility at least two previous generations (I.e., 8GFC backward compatible to 4GFC and 2GFC.*
- *†Line Rate: All "FC" speeds are single-lane serial stream*
- *‡Dates: Future dates estimated*

Continued...FCIA Maps FC to FCoE

ISL (Inter-Switch Link) Roadmap

Product Naming	Throughput (MBps)	Equivalent Line Rate (GBaud)†	T11 Spec Technically Completed (Year) ‡	Market Availability (Year) ‡
10GFC	2400	10.52	2003	2004
20GFC	4800	21.04	TBD	2008
40GFC/ FCoE	9600	41.225	2010	MARKET DEMAND
100GFC/ FCoE	24000	103.125	2010	MARKET DEMAND
400GFC/ FCoE	96000	TBD	TBD	MARKET DEMAND
1TFC/FCoE	240000	TBD	TBD	MARKET DEMAND

- *ISLs are used for non-edge, core connections, and other high speed applications demanding maximum bandwidth. Except for 100GFC (which follow Ethernet)*
- *†Equivalent Line Rate: Rates listed are equivalent data rates for serial stream methodologies.*
- *‡ Some solutions are Pre-Standard Solutions: There are several methods used in the industry to aggregate and/or "trunk" 2 or more ports and/or data stream lines to achieve the core bandwidth necessary for the application. Some solutions follow Ethernet standards and compatibility guidelines. Refer to the FCoE page 4 for 40GFCoE and 100GFCoE.*

FCoE Roadmap

Product Naming	Throughput (MBps)	Equivalent Line Rate (GBaud)†	T11 Spec Technically Completed (Year) ‡	Market Availability (Year) ‡
10GFCoE	2400	10.3125	2008	2009
40GFCoE	9600	41.225	2010*	MARKET DEMAND
100GFCoE	24000	103.125	2010*	MARKET DEMAND

- *Fibre Channel over Ethernet tunnels FC through Ethernet. For compatibility all 10GFCoE FCFs and CNAs are expected to use SFP+ devices, allowing the use of all standard and non standard optical technologies and additionally allowing the use of direct connect cables using the SFP+ electrical interface. FCoE ports otherwise follow Ethernet standards and compatibility guidelines.*
- *†Line Rate: All "FC" speeds are single-lane serial stream*
- *‡Dates: Future dates estimated*
- **It is expected that 40GFCoE and 100GFCoE based on 2010 standards will be used exclusively for Inter-Switch Link cores, thereby maintaining 10GFCoE as the predominant FCoE edge connection*

TWICE AS FAST IS MORE THAN TWO TIMES BETTER

Scott Kipp – Senior Technologist, Brocade

Fibre Channel moves to a new speed to help users

Fibre Channel, the dominant storage networking technology, is developing a new round of products to reinvigorate Fibre Channel. 16 Gigabit/second Fibre Channel (16GFC) products will be released in 2011 and a loyal base of users will migrate to the next generation of speed. 16GFC is backward compatible with earlier Fibre Channel speeds and will continue to use SFP+ optical modules so that Fibre Channel has economies of scale unlike any other multi-Gigabit/second technology.

While Fibre Channel is often considered a niche market, the facts speak differently for multi-gigabit technologies. According to Dell'Oro, 10 Gigabit Ethernet sold about 2.1 million (2.1M) switch ports in 2009 while Fibre Channel sold 5.8M switch ports. The physical layer of Fibre Channel has more economies of scale than 10GbE because Fibre Channel only uses the SFP+ form factor while 10GbE uses the XENPAK, X2, XFP and SFP+ form factors. Over 95% of Fibre Channel links are within the data center over multimode fiber while 10GbE operates over a variety of distances and single-mode fiber to fracture the 10GbE physical market into multiple pieces. While 3.5M 8GFC switch ports shipped in 2009, Dell'Oro forecasts 5.7M 8GFC switch ports in 2010. The niche technology of Fibre Channel is growing by leaps and bounds and set to leapfrog 10GbE technologies in terms of cost, speed and capabilities.

With the release of a new generation of 16GFC ASICs and products, Fibre Channel will continue its dominance in storage networking. 16GFC doubles the data throughput of 8GFC links from 800 MegaBytes/second (MBps) to 1,600 MBps with 16GFC. From Host Bus Adapters (HBA) to switches, 16GFC will enable higher performance with lower power consumption per bit over previous generations of technology. 16GFC delivers the performance required by today's leading applications.

The benefits of any 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 16GFC is used instead of 10GbE, 8GFC or 4GFC. Several technology advances are pushing up bandwidth demands in SANs that include application growth, server virtualization, multi-core processors, PCI Express 3.0, increased memory and solid state disks. 16GFC is keeping pace with other technology advances in the data center.

16GFC should be applied where high bandwidth is needed. Applications where 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. 16GFC is designed to assist users in transferring large amounts of data and decreasing the number of links in the data center.

Overview of 16GFC

16GFC has considerable improvements from the previous Fibre Channel speeds that include using 64b/66b encoding, transmitter training and linear variants as outlined in Table 1. 16GFC uses electronic dispersion compensation (EDC) and transmitter training to improve backplane links. The combination of these technologies enables 16GFC to provide the highest throughput density in the industry.

Continued...Twice as Fast is More Than Two Times Better

To remain backward compatible with previous Fibre Channel speeds, 16GFC ASICs must support 8GFC and 4GFC to meet the Fibre Channel Industry Association's roadmap for future speeds. The 16GFC ASICs must have 8b/10b codecs for 4GFC and 8GFC and 64b/66b codecs for 16GFC. Users can attach new 16GFC devices and switches to existing infrastructure and the 16GFC devices will auto-negotiate down to the lower speeds of the legacy devices. The new 16GFC ports can be seamlessly added to the existing networks to increase performance for new segments of the storage network while not requiring a forklift upgrade. Fibre Channel has proven that it can jump to new speeds effectively and 16GFC is a continuation of that long range thinking that has made Fibre Channel such a success.

Fibre Channel Speed Characteristics

Speed Name	Throughput (MBps)	Line Rate (Gbps)	Encoding	Retimers In The Module	Transmitter Training
1GFC	100	1.0625	8b/10b	No	No
2GFC	200	2.125	8b/10b	No	No
4GFC	400	4.25	8b/10b	No	No
8GFC	800	8.5	8b/10b	No	No
10GFC	1200	10.53	64b/66b	Yes	No
16GFC	1600	14.025	64b/66b	Yes	Yes

Table 1: Fibre Channel Speed Characteristics

The Benefits of Higher Speed

The benefits of faster tools are always the same – more work in less time. By doubling the speed, 16GFC 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 16GFC 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

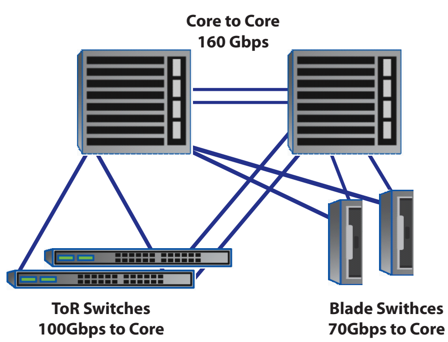

Figure 1: Network Design Implications

	8GFC Links	16GFC Links
ISLs from ToR Switch to Core	16	8
ISL from Blade Swithc to Core	10	5
Core to Core	24	12
Total ISLs	50	25
Total Ports	100	50

Besides the reduction in equipment that cuts power consumption dramatically, 16GFC 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 16GFC designs is for a 16GFC port to consume less power than two 8GFC links that deliver the same throughput. Initial estimates for power consumption show 16GFC SFP+s consuming 0.75 Watts of power while 8GFC SFP+ consuming 0.5 Watts of power. These estimates show that a 16GFC link will consume 25% less power than two 8GFC ports.

Continued...Twice as Fast is More Than Two Times Better

If fewer links are needed, cable management becomes simpler. Managing cables behind a desktop or home entertainment center are bad enough, but managing hundreds of cables from a single switch or bundles of cable from a server can be horrendous. The reduction of cables aids in 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 16GFC links aids cable management.

The end result of 16GFC is that there are less links, less cables, less ports and less power for the same performance. Figure 2 shows the comparison of one 16GFC link to two 8GFC links. The largest benefits of the 16GFC ports will be the fewer number of HBAs and switch ports that are connected to these media.

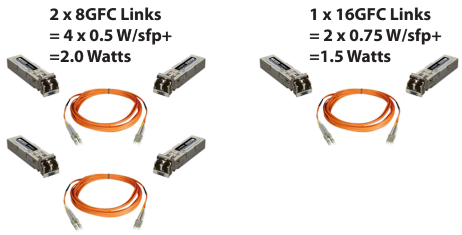

Figure 2: Comparison of media

Applications of 16GFC

16GFC is designed for high bandwidth applications that include:

- ISLs – covered in previous section
- High Bandwidth Applications
- Data Migration
- Virtual Desktop Infrastructure (VDI)
- Solid State Disks or memory arrays

The type of servers that need 16GFC HBAs will 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, the servers need to be designed for the demanding times like backup and data mining (analytics) when 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 16GFC is streaming I/O applications. 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 and controllers and drives are sending out sequential reads or writes as fast as they can.

Another application of 16GFC links in the data center is between data centers, storage arrays or clouds. During data center consolidations, disaster recovery and equipment changes, users often have a need to migrate terabytes (TBs) or even petabytes (PBs) 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. Table 2 shows the time required to transfer large amounts of data at 1,600 MBps with 16GFC. When time is money, 16GFC is better.

Data Migration

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 Wekk

Table2: Data Migration Examples

VDI is a growing trend in enterprises where virtual desktops in the data center are sent to users on a variety of devices. VDI has the advantage of centralized management where applications and hardware can be easily upgraded 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. 16GFC is one of the many components that can help improve performance at these critical initialization times.

Storage arrays based on memory or solid state disks (SSDs) are enabling a new level of performance in high performance computing. With lower latency and higher IOPs than traditional storage arrays, 16GFC 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 applications like online gaming where these applications have already reached bandwidth requirements of 50 GB/s. With the price of SSDs dropping quickly, SSDs should be able to address many more applications where performance is more important than capacity.

SUMMARY

Speed wins! It's not rocket science to understand that a link that is twice as fast as a slower link can do more work. While many applications won't use the full extent of a 16GFC link yet, over the next few years, traffic and applications will grow to fill the capacity of 16GFC. The refresh cycle for networks is often longer than that of servers and storage, so 16GFC 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 8GFC. To future-proof deployments, 16GFC should be considered to be the most efficient way to transfer large amounts of data in data centers. With proprietary trunking technology at 16GFC, users can get up to 128GFC of performance that delivers more bandwidth per power and cost.

16GFC will be the best performer in several applications. 16GFC can reduce the number of ISLs in the data center or migrates a large amount of data for array migration or disaster recovery. High performance applications like VDI that use SSDs or require high bandwidth are ideal applications for 16GFC. 16GFC combines the latest technologies in an energy efficient manner to provide the highest performing SANs in the world. With these benefits, 16GFC is more than two times better than 8GFC.

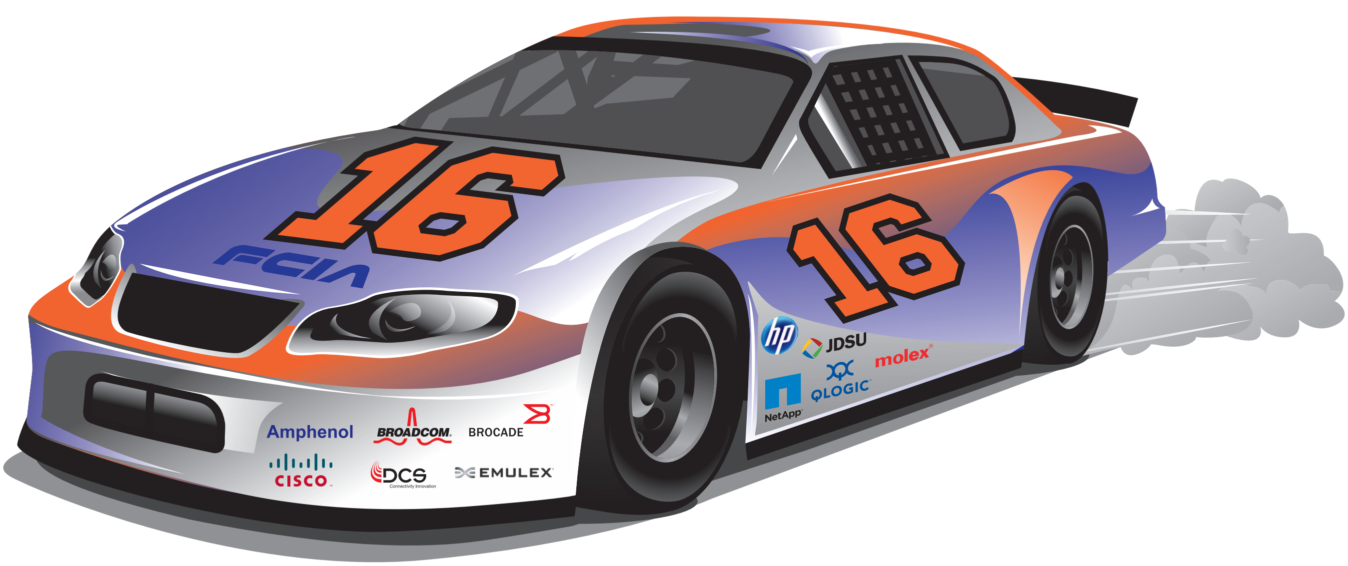

UNIFIED FABRIC AND FCoE

J Metz, Program Manager FCoE, Cisco Systems

Introduction

The question is actually quite simple: "What vision do you have for the Data Center?" The answer, of course, is never as simple (or easy) as the question.

One of the characteristics of data centers is that they happen to be unique. Like fingerprints, the sheer number of options for servers, networking and storage means that there are often disparate groups of technology and capabilities, often leading to underutilized resources in some areas, and oversubscribed resources in others.

Most IT Directors would like to eliminate the former and mitigate the latter.

The IT industry is on a journey to move workloads from dedicated physical resources, to dynamic virtual resources, to automatically scaling resources. Businesses demand application delivery to be seamless and instantaneous, which in turn requires an infrastructure that rapidly scales and adjust to variable (yet increasing) workload demands

This requires a certain degree of elasticity and agility. More and more often the IT platforms that can increase capacity and reduce power as needed in a dynamic response must be done with more common equipment and infrastructure. When looking forward into the evolution of data centers, the key resounding principle that seems to resonate appears to be an ability to combine, converge, or unify similar functionality where possible. Thus, the term "unified" was born.

As noted elsewhere in this solution guide (see "Unified Fabric Management"), the terms "Unified" and "Converged" can mean different things to different people (or vendors!) and can often lead to confusion. In recent months it can likely lead to audience fatigue, as "just another" marketing blitz.

Nevertheless, the consequences of this phenomenon are very real and have long-term implications long after the marketing buzz has moved on to another term. No matter what the name, these issues still need to be addressed.

What does Unified/Convergence mean?

It is generally understood that one of the main culprits of this inefficiency lies within the technology silos: networking, servers, storage and application servers. Herein lies the challenge to reduce the unnecessary duplication and organizational waste that naturally results from suboptimal resource utilization. Obviously if there is a way to combine siloed technologies to perform better, there is a compelling reason to explore these alternatives further.

However, up until recently there have been limitations to doing 'convergence' within a data center. Most notably – especially within the storage networking space – the ability to use a single, ubiquitous networks risk storage best practices.

Therein lies the problem, though. How do you consolidate equipment to eliminate potential waste while still retaining the best practices and design principles required for high availability, reliability, and failover scenarios?

To that end, then, a unified/converged environment must preserve the best practices of both LAN and SAN designs. Otherwise, if that is not the case, if one is subordinated to the other, this is not "unification" or "convergence:" it's annexing!

Continued...Unified Fabric and FCoE

In this way, storage administrators rightfully get concerned when it comes to talk of convergence or unification, because they have spent their entire careers mitigating and reducing risk. The risk of the unknown is a powerful obstacle for assuaging the fears of those who are responsible and accountable for data integrity.

What role does FCoE play?

Isolating the discussion to just LAN and SAN environments, we often find ourselves at diametrically opposite philosophies on best practices. Take a look at Table 1, for example, and see where some of the inconsistencies and incompatibilities are the most obvious:

Characteristic	LAN	WAN
Client-Server Relationships	Not Pre-Defined	Pre-Defined
Transport	Lossy	Lossless
Services	Relies on upper layers, often in different devices	Same layer as transport in the same devices
High Availability	Service HA implemented separately	Provided through dual-fabric architecture

How can these major differences be reconciled? These differences exist all the way down to the wire, so how can we find a single technology to solve the problem? This has been a major obstacle to overcome when attempting to collapse the silos into a single ubiquitous environment.

One of the downfalls of any new technology that is introduced into the marketplace is the inevitable comparison as a replacement for existing technologies. The claim that iSCSI is a "Fibre Channel-Killer" never came to fruition because the market found room for both technologies within the Data Center at the same time. In short, iSCSI became another tool in the Storage administrator's toolbox.

FCoE, however, relies on enhancements to Ethernet that allows storage administrators to maintain these characteristics on the same wire as LAN systems that have different requirements. Because FCoE is Fibre Channel, it allows us to maintain the pre-defined relationships between clients and servers, and permits us all the zoning services in the same devices that SANs require. Not only that, simultaneous, dual-fabric architectures remain the best practice for FCoE-based SANs.

The underlying Ethernet connection, approved and standardized in its final form in 2011 by IEEE, makes several technologies available for implementation that makes converged networks possible:

- 802.1p provides a mechanism for implementing Quality of Service at the MAC level, which in turn allows different types of traffic to operate on the same wire with no covariance on individual priorities (channels).
- 802.1Qbb – Priority Flow Control (PFC) provides the same functionality (although with a different mechanism) as Fibre Channel buffer-to-buffer credits to control the traffic on the correct priority. In short, this is what makes the FCoE traffic lossless over Ethernet.
- 802.1Qaz – Enhanced Transmission Selection (ETS) and Data Center Bridging eXchange (DCBX) round out our new technologies by permitting the finer-grained tuning of bandwidth when links are shared, as well as provides the hand-shaking and configuration validation between devices.

In other words, by being able to segment out the traffic on a 10 Gb Ethernet link and apply behavioral characteristics to each independently, it is possible to consolidate both LAN and SAN traffic types (or unify, or converge the networks) safely and reliably.

The Big Picture

It cannot be understated to stress that without these details the big picture of a Unified Fabric is unattainable. Without the ability to provide traffic-engineering capabilities unique to SANs while simultaneously running LAN topologies, broader concepts such as topology or architecture considerations would be moot.

Continued...Unified Fabric and FCoE

Traditional LAN architectures can not maintain Fibre Channel's best practices, high availability, high performance, or troubleshooting mechanisms. In a converged network, however, customer implementations open up considerably with a wide variety of deployment choices.

What choices? The choices that come from being able to run any protocol at any time on any piece of equipment on an as-needed basis without having to worry about the underlying infrastructure. Have a need for iSCSI? You have the infrastructure ready to use iSCSI priorities. Need Fibre Channel protocol for deterministic traffic for a Tier-1 application? You have the infrastructure ready for that.

There's also the elimination of "budget-" and "bandwidth-leapfrogging." Because of the way that Data Centers have been siloed into various categories, budget considerations and lifecycles are often uncoordinated and independent from each other. Often the dependencies upon which oversubscription ratios are calculated are "best effort" or "best guess" attempts, especially given the tendency to attempt to plan today for the needs 4 or 5 years down the road.

It is rare that these needs are in harmonious sync with each other.

Should an application need to be rolled out without the ability to provision the capacity, it is not possible to simply 'borrow' from the other technologies. For instance, we cannot simply borrow the extra underutilized capacity of a Director-Class Ethernet switch and give it to a Fibre Channel SAN using traditional, unconsolidated architectures, or vice versa.

With a converged network, however, the issue is one of configuration and capacity allocation, of re-allocating ports and bandwidth to accommodate the new demands without impacting the existing physical infrastructure. Instead of managing multiple network capacity roadmaps there is only one, unified fabric. This simplifies (and reduces) considerable OpEx expenses over time, not just the initial CapEx outlays for only needing to purchase a single type of network: the converged type.

Conclusion

Despite the issue surrounding how the terms are used, or the risk of them becoming overexposed, the needs of the evolving Data Center are very real and can have profound long-term consequences.

The impact of Unified Fabrics and Converged Networks are only in their embryonic stages, and as IT departments begin to learn more about the technologies there will be greater and unanticipated uses for the technology that will change the way that we think of providing network-based storage in new and exciting ways. FCoE is only one part of this, but a valuable cornerstone of ensuring a solid, reliable and predictable means for getting the best of both worlds.

UNIFIED FABRIC MANAGEMENT

Rupin Mohan, Business & Technology Strategy Manager, Hewlet-Packard, and J Metz, Program Manager FCoE, Cisco Systems

"Unified' and 'Converged' may be the most over used terms in technology now days. In this age of austerity, it has become the trend to combine two or more technology functions, be it software or hardware and call it either converged or unified.

Let's get more specific to data center networking. With the development of Fibre Channel over Ethernet (FCoE), customers have a choice to run IP traffic (LAN) and storage (Fibre Channel SAN) traffic on the same physical wire and switches. If you are using multiple Ethernet and Fibre Channel connections and wires to manage the network and storage traffic in your operating environment, then you are likely concerned about the density and sprawl of switches in your data center and related management and maintenance costs to effectively manage this network. As a potential solution to the sprawl and increased management costs, FCoE as a technology should be on your radar.

As Data Center architectures become more and more capable, the need for comprehensive management tools are more and more pressing. In fact, they're less of a luxury and more of a requirement. But, the sixty four thousand dollar question is - what qualifies as a "comprehensive management tool?" Vendor implementations take on various metaphors for what a "unified fabric" is versus what "converged fabric" means.

For instance, one view refers to 'unified' as the large-scale symbiosis of servers, storage, and network, while 'converged' or 'flex' refers to the ability to transport different types of LAN and SAN traffic across the same networking infrastructure. Another popular view is exactly the opposite where 'converged' refers to servers, storage and network in a box and 'unified' or 'flex' refers to the ability to transport different types of LAN and SAN traffic across the same networking infrastructure. Still others may have their own mechanisms for defining terminology.

Rather than choose sides or be proscriptive, this article examines the criteria that should be used to determine the appropriate tools that can help you identify what unification management looks like.

There are several other dimensions to this story. Let's discuss three:

1. **Rise of solutions comprising of server, network (IP and/or FC) and storage in a box.**
   These are sometimes tightly and sometimes loosely integrated. But, overall, the big value proposition is that these solutions are fully qualified, easy to buy, easy to deploy and easy to service from a customer point of view. This is what the industry is generally referring to as 'Converged Infrastructure' or "Unified Computing'. A flexible network / fabric is one of the key architectural tenets of this solution.
2. **Server virtualization gaining significant traction in the market.**
   When customers think of Cloud computing, the first thing that comes to their mind is server virtualization. If you think about it, server virtualization serves as the prime example of unification / convergence in today's data center. In this age of more of 'more for less', convergence from server virtualization delivers tangible results. Note that multi-core CPU's, dramatic increases in server I/O bandwidth, and high speed network interconnects like 8Gb Fibre Channel and 10GbE DCB Ethernet have been the key foundational drivers of server virtualization.

Continued...Unified Fabric Management

3. **Management Software becoming a control point or sticking point in the data center.** The rise of management software from server virtualization vendors as a 'unified' management tool which manages your server, network and storage assets is a reality. There is also software from switching providers attempting to manage unified storage area networks and Ethernet networks. This is an area again where there are products that are loosely coupled but strategically, vendors are talking about a single software management platform that manages the 'unified' or 'converged' network in the data center.

Let's dig a little deeper into management software. Every storage and storage network vendor provides some level of management for its products. The issue really begins and ends with the concept of scale, and what happens when many pieces of equipment are connected together – especially in a heterogeneous environment.

To that end, the 5 basic things you need to look out for are:

Comprehensive, ground-up approach to end-to-end management
Single versus heterogeneous support
Role-based access controls
Threshold alerts and Trending analysis
3rd Party API integration

Let's take a look at these in turn.

1. Comprehensive, ground-up approach to converged/unified network management

As most administrators are painfully aware, network and SAN equipment providers historically have released loosely coupled management tools primarily to be able to check that box in the RFP's. This often leads to having different management parameters and requirements across the same vendor portfolio. This can pose some issues, particularly in the area of unintended consequences between the interconnection and interrelationship between pieces of equipment.

Understanding the traffic within a network takes more than just managing at a device-level. When it comes to storage it is critically important to understand from the "ground up," from the frame's initiation point at the initiator, all the way through to its target. Requirements for IP traffic are different compared to requirements of storage traffic in terms of latencies, packet drop ratios, timeouts, etc. This means being able to track the elements where frames may run into potential problems becomes more important. Understanding where problems may occur and make it visibly easy to identify at a glance become extremely useful in preventing issues before they become critical.

Network management tools should allow you the ability to use pre-existing 'canned' templates or customized configurations that can help rapid deployment of new equipment with pre-set settings. Wizard-based functionality may also be available for automating mundane tasks that may be error-prone (e.g., such as device zoning, converged LAN/SAN port configuration, etc.).

2. Single versus heterogeneous support

There is another interesting angle that could be important to data center administers and that is if this software manages products from a single vendor versus products from heterogeneous vendors. Customers are realizing the value of competition and dual sourcing when it comes to purchasing IT assets and also of buying best in class products. Hence, the importance of heterogeneous support is increasing and any management software that can manage all these assets from a single pane of glass does add considerable value to the administrator.

3. Role-based access controls

Whether you call the environment a "unified fabric" or "converged network," one thing remains the same: eventually there will be multiple teams of people interacting with the same equipment. From a storage perspective, the obvious questions arise: How do you prevent someone from messing up your storage, whether intentionally or by accident?

Continued...Unified Fabric Management

Without question the ability to set access control limits based upon roles is an unqualified "must." Any time there is the potential for conflict when using ubiquitous equipment there must be a means to mitigate it within software, regardless of how the network/storage organizational structure works within your company, the management tool should provide this functionality.

4. Threshold alerts and trending analysis/reports

Performance monitoring is a key element to end-to-end network management, whether it be for LAN or SAN environments. Management software should – at a minimum – be able to identify the current and existing traffic patterns from source to destination for all types of traffic, and be able to easily and quickly distinguish between them.

In addition, there should be an ability to track and analyze network performance over time. By being able to run traffic prediction analyses these tools can often enable the administrator to run "what if" scenarios within the data center, including workflow mobility impacts and virtualization loads on storage arrays, not to mention the impact of adding new applications to the network load.

Logging and reporting becomes critical in these cases. Some organizations wish to begin chargeback to departments for these services, particularly as companies begin to examine the notion of "private clouds" with Infrastructure-as-a-Service (IaaS).

5. 3rd Party API Integration

Another interesting development in the management software area is the rising importance and interest in developing 3'rd party API's to the core management modules. The core idea behind this is similar to the rise of applications in smart phones. Vendors are expecting customers and OEMs to develop custom applications using these 3'rd party API's which solve individual problems and add specific value. Whether it be a storage tool that must connect with a server virtualization tool, or a networking tool that needs to have interactivity with servers, management tools must at least be able to call on each others' APIs to be able to poll and process critical data. These API's further contribute to the 'network effect' in the data center network wherein customers get hooked on to features and functionality from a particular vendor.

Conclusion

As you can see, this is an area of enormous innovation and intense competition where vendors are working hard to differentiate and create value for customers. The taxonomy is evolving and requirements for unified / converged management software are being developing. One thing is certain, that management software would be a significant control point for in the data center for converged / unified solutions.

2010/2011 FCIA Board of Directors

FCIA Officers

- Skip Jones, QLogic Corporation: Chairman
- Mark Jones, Emulex: Vice Chairman
- Greg McSorely, Amphenol: Treasurer
- Jay Neer, Molex: Secretary

Members at Large

- Anthony Faustini, Cisco Systems
- Kevin Ehringer, DCS
- Rupin Mohan, Hewlett-Packard
- Rick Jooss, NetApp
- Steven Wilson, Brocade
- Eisaku Saili, Netmarks Inc., FCIA Japan BOD Liaison

FCIA Solutions Guide Acknowledgements

- Michael LoBue - Marketing Manager, FCIA
- Chris Lyon - Executive Director, FCIA
- J Metz, Program Manager FCoE, Cisco Systems
- Mike Tenhulzen - Marketing Consultant
- Rupin Mohan - Business & Technology Strategy Manager, Hewlet-Packard
- Scott Kipp - Senior Technologist, Brocade
- Skip Jones - Director of Technology and Planning, QLogic: FCIA Chairman
- Steve Wilson - Director, Technology and Standards, Brocade: INCITS TC T11 Chairman

FCIA
FIBRE CHANNEL INDUSTRY ASSOCIATION

About the FCIA

The Fibre Channel Industry Association (FCIA) is a non-profit international organization whose sole purpose is to act as the independent technology and marketing voice of the Fibre Channel industry.

We are committed to helping member organizations promote and position Fibre Channel, and to providing a focal point for Fibre Channel information, standards advocacy, and education.

Today, Fibre Channel technology continues to be the data center standard for storage area networks and enterprise storage, with more than 80 percent market share.

Join FCIA

Learn more about member benefits, download a membership application, read the by-laws, on our website at:
http://www.fibrechannel.org/FCIA/appbylaws.html

Contact FCIA

For more information, go to our website at:

- www.fibrechannel.org
- email: info@fibrechannel.org

FCIA

FIBRE CHANNEL INDUSTRY ASSOCIATION

TECHNICAL AND MARKETING VOICE OF THE FIBRE CHANNEL INDUSTRY

fibrechannel.org