



# FIBRE CHANNEL

SOLUTION GUIDE 2012 - 2013



[FIBRECHANNEL.ORG](http://FIBRECHANNEL.ORG)



FIBRE CHANNEL INDUSTRY ASSOCIATION

A vertical photograph on the left side of the page shows light trails from a city at night. The trails are in shades of blue, cyan, red, and white, creating a sense of motion and energy. In the background, a building is visible with some lights on.

# FIBRE CHANNEL AND FCoE

POWERING THE NEXT GENERATION  
PRIVATE, PUBLIC, AND HYBRID  
CLOUD STORAGE NETWORKS

## ➤ ABOUT THE FCIA

The Fibre Channel Industry Association (FCIA) is a non-profit international organization whose sole purpose is to be 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.

## ➤ CONTACT THE FCIA

For more information:  
[www.fibrechannel.org](http://www.fibrechannel.org)  
[info@fibrechannel.org](mailto:info@fibrechannel.org)

## ➤ CONTENT

- 4 STATE OF THE FIBRE CHANNEL INDUSTRY
- 6 SPEED AND CONVERGENCE: FOLLOW THE FCIA ROADMAP TO SUCCESS
- 9 16GFC DEPLOYMENTS ADDRESSES DATA CENTER CHALLENGES
- 13 40Gb FCoE: THE FUTURE OF HYPER-AGILITY IN THE DATA CENTER
- 16 FIBRE CHANNEL CONNECTIVITY
- 18 HOW FIBRE CHANNEL PLUGFESTS HELP PREPARE PRODUCTS FOR MARKET

# ➤ STATE OF FIBRE CHANNEL INDUSTRY

This year, Fibre Channel has achieved a significant milestone. FCIA has termed it 10-10-10 for 10 million Fibre Channel ports shipped, \$10B invested in Fibre Channel technology and 10 Exabytes (EB) of storage shipped . In 2012, 10M Fibre Channel ports are expected to ship when you combine adapter and switch ports<sup>[1]</sup>, not even counting target ports on storage arrays. Also, over \$10B of Fibre Channel Enterprise Storage Systems will be sold worldwide, representing 62% of all Storage Area Network storage in the market<sup>[2]</sup>. The FCIA calculates that over \$100B has been invested in Fibre Channel over the last two decades. Finally, for the first time ever, 10 EB of Fibre Channel storage capacity in external enterprise storage systems will be shipped - representing 66% of the SAN market<sup>[3]</sup>. So, 2012 is the year of 10-10-10 for Fibre Channel that shows how Fibre Channel continues to flourish with the rest of the IT markets.



10 - 10 - 10

[1] Worldwide Storage Area Network Market – Fibre Channel Forecast, January 2012

[2] Worldwide External Enterprise Storage Systems Revenue by Topology, Installation, and Protocol 2006-2015 (\$B), IDC, 2011

[3] Worldwide External Enterprise Storage Systems Capacity Shipped by Topology, Installation and Protocol 2006-2015 (PB), IDC, 2011

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 two decades, 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 to 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 1GFC, 2GFC, 4GFC, 8GFC and 16GFC communications, with 32GFC industry standards development complete this year, 2012, and 10GFCoE, 40GFCoE and 100GFCoE providing an ideal solution for fast, reliable mission-critical information storage and retrieval for today's data centers.

**10 MILLION FIBRE CHANNEL PORTS**

**\$10B INVESTED IN FIBRE CHANNEL TECHNOLOGY**

**10 EXABYTES (EB) OF STORAGE SHIPPED**

If we look at what is coming in 2012 (Virtualization, Big Data and Analysis, Cloud storage networks, social media, archive, disaster recovery, personal, professional and corporate media, records, entertainment, and in general e-accessibility to most all things, IO bandwidth requirements are greatly increasing and going nowhere but up. IO bandwidth will be expected to handle more and more traffic, be it the same dedicated traffic types on separate physical networks, or multiple converged data traffic types onto a single network.

# ▶ SPEED AND CONVERGENCE

FOLLOWING THE FCIA ROADMAP TO SUCCESS!



By: Skip Jones - Chairman, FCIA

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 of a technology. It is also a guide to 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.

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 for the future.

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.

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 with 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. It takes time 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 did not frequently change and reset expectations in the industry. A changing roadmap causes confusion and could cause faulty planning by users and suppliers 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 15 years of accurate roadmaps come from the seriousness FCIA takes in this huge responsibility and obligation to the industry.

## FIBRE CHANNEL ROADMAP

- “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

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

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 includes 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 both feasible within the roadmap timelines, and they also result 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.

## ISL (INTER-SWITCH LINK) ROADMAP

- 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 roadmap for 40GFCoE and 100GFCoE.

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

## FOLLOWING THE FCIA ROADMAP TO SUCCESS!

FCIA's process for 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 its guidelines.

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 demand.

FCIA's Roadmap, version V14, is the latest descendent of a long successful history of the FCIA roadmap and can be found at: [www.fibrechannel.org/fibre-channel-roadmaps.html](http://www.fibrechannel.org/fibre-channel-roadmaps.html). 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 V14 Roadmap also pinpoints standard stability and general market availability for 16GFC and 32GFC edge connectivity (16GFC in 2011 and 32GFC in 2014, respectively). 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.

### FCoE ROADMAP

- 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

Product Naming	Throughput (MBps)	Equivalent Line Rate (GBaud)†	T11 Spec Technically Completed (Year) ‡	Market Availability (Year) ‡
10GFCoE	2400	10.315	2008	2009
40GFCoE	9600	41.225	2010*	Market Demand
100GFCoE	240000	103.125	2010*	Market Demand



# 16GFC DEPLOYMENT ADDRESSES DATA CENTER CHALLENGES

16 GIGABIT PER SECOND FIBRE CHANNEL (16GFC) PRODUCTS WERE RELEASED IN 2011 AND ARE BEING WIDELY DEPLOYED.

By: Scott Kipp, Senior Technologist, Brocade

Rupin Mohan, Senior Manager Strategy & Product Planning, Hewlett Packward and Member of FCIA Board of Directors

## Introduction

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

16GFC is backward-compatible with 3 previous generations of Fibre channel – 2GFC, 4GFC and 8GFC so that customers can leverage their existing Fibre Channel investments in their data center. Using the same Small Form Factor Pluggable + (SFP+) interface, 16GFC enables the use of the same cabling infrastructure that they have been deploying for years.

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 8GFC or 4GFC. Several technology advances in the data center 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 server virtualization, 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 technical improvements from the previous Fibre Channel speeds that include using 64b/66b encoding, transmitter training and linear variants as outlined in Table 1. 16GFC doubles the throughput of 8GFC to 1,600 MBps but uses 64b/66b encoding to increase the efficiency of the link. 16GFC links also use retimers in the optical modules to improve link performance characteristics. 16GFC also 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.

Speed Name	Throughput (MBps)	Line Rate (GBaud)	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

While 16GFC doubles the throughput of 8GFC to 1600 MBps, the line rate of the signals only increases to 14.025 Gbps because of a more efficient encoding scheme. Like 10GFC and 10 Gigabit Ethernet, 16GFC 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 16GFC, 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 off the line rate so that the links could run over 100 meters of optical multimode 3 (OM3) fiber. By using 64b/66b encoding, 16GFC improves the performance of the link with minimal increase in cost.

To remain backward compatible with previous Fibre Channel speeds, Fibre Channel application specific integrated circuit (ASIC) must support 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 towards 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.

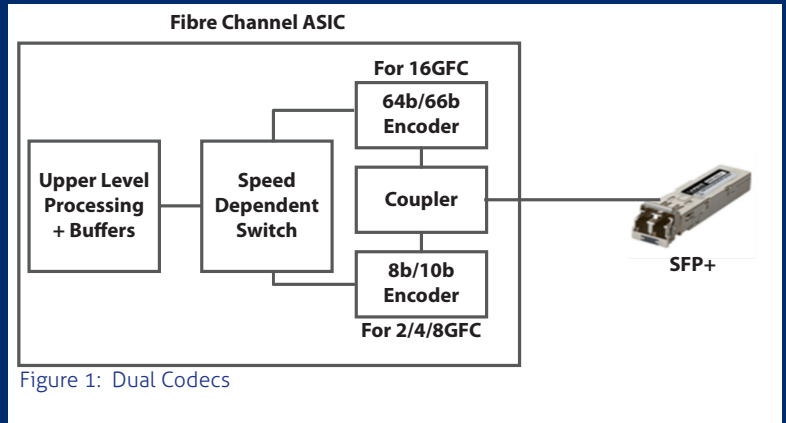


Figure 1: Dual Codecs

The second technique that 16GFC 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 the majority of the jitter in a link. By placing a retimer in the optical modules, the link characteristics are improved so that the links can be extended for optical fiber distances of 100 meters on optical multimode 3 (OM3) fiber. The cost and size of retimers has decreased significantly so that they can be integrated into the modules for minimal cost.

The 16GFC 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. 16GFC 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 16GFC can support 125 meters. If a 16GFC 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 16GFC to work in a wide range of applications.

Speed Name	Multimode OM1 Link Distance	Multimode OM2 Link Distance	Multimode OM3 Link Distance	Multimode OM4 Link Distance	Single-mode OS1 Link Distance
	62.5 um core and 200 MHz*km	50 um core and 500 MHz*km	50um core and 2000 MHz*km	50um core and 4700 MHz*km	9um core and ~infinite MHz*km
1GFC	300	500	860	*	10,000
2GFC	150	300	500	*	10,000
4GFC	50	150	380	400	10,000
8GFC	21	50	150	190	10,000
10GFC	33	82	300	*	10,000
16GFC	15	35	100	125	10,000

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

Table 2: Link Distance with Speed and Fiber Type (meters)

Another important feature of 16GFC is that it uses transmitter training for backplane links. Transmitter training is an interactive process between the electrical transmitter and receiver that tune lanes for optimal performance. 16GFC 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 16GFC 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, 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

### Reduced Number of Links

As with other speeds of Fibre Channel, 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 16GFC to yield a 128GFC link between any two switches. These trunks can grow from 16GFC to 128GFC in 16G increments.

Figure 2 shows a simple comparison of the number of links in an 8GFC fabric and a 16GFC fabric. The higher speed links of 16GFC eliminates tens or hundreds of ports from a comparable 8GFC fabric. The real savings occur when the number of HBAs, switches and end devices can be decreased with the higher performance of 16GFC. In the example in Figure 2, a Top of Rack (ToR) switch needs 100 Gbps of bandwidth so the user needs eight 16GFC ISLs instead of sixteen 8GFC ISLs. Similar comparisons between 16GFC ISLs and 8GFC ISLs are given in the table in Figure 2 to show how fewer ports and links are needed at 16GFC.



	8GFC Links	16GFC Links
ISLs from ToR Switch to Core	16	8
ISL from Blade Switch to Core	20	5
Core to Core	24	12
Total ISLs	50	25
Total Ports	100	50

Figure 2: Network Design Implications

### Reduced Power Consumption per Bit

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.

### Easier Cable Management

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.



**SPEED  
WINS!**

The Fibre Channel industry has teamed with cable management companies to provide incredibly dense solutions for cable management. To reduce the bulk of cables, multiple vendors offer uniboot cables that combine two fibers into one cord and then combining 12-fibers into one ribbon cable as shown in Figure 3. The LC to MPO cable harnesses reduce cable bulk and utilize compact fiber ribbons. The fiber optics cable industry also provides very dense patch panels for MTP and LC connectors - excellent solutions for all varieties of cable management.



Figure 3: Uniboot LC to MTP Cable Harness

### Summary of Benefits

The end result of 16GFC is that there are less links, less cables, less ports and less power for the same performance. Figure 4 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 4: Comparison of media

### 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. Switch vendors with trunking technology can scale to 128Gbps of performance between two points. 16GFC will be the best performer in several applications. 16GFC 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 like virtual desktop infrastructure and solid state disks (SSDs) require high bandwidth are ideal applications for 16GFC. As more applications demand the low-latency performance of SSDs, 16GFC keeps up with other advances in other components of the storage infrastructure. 16GFC combines the latest technologies in an energy efficient manner to provide the highest performing SANs in the world.

# > 40Gb FCoE

THE FUTURE OF HYPER-AGILITY IN THE DATA CENTER

By: J Metz, Ph.D., Cisco, Product Manager – Storage

## INTRODUCTION

It may sound strange to think of the Fibre Channel Industry Association discussing Ethernet technologies in a Fibre Channel solution guide. After all, when people think of Fibre Channel something more than just the protocol comes to mind - the entire ecosystem, management, and design philosophies are part and parcel of what storage administrators think of when we discuss “Fibre Channel networks.”

There is a reason for this. Over the years, Fibre Channel (FC) has proven itself to be the benchmark standard for storage networks – providing well-defined rules, unmatched performance and scalability, as well as rock solid reliability. In fact, it’s a testament to the foresight and planning of the T11 technical committee that the FC protocol is robust enough to be used in a variety of ways, and over a variety of media.

Did you know, for instance, that the T11 committee has created a number of possible forms for transporting Fibre Channel frames? In addition to the Fibre Channel physical layer, you can also run FC over (though not an exhaustive list):

- Data Center Ethernet
- TCP/IP
- Multiprotocol Label Switching (MPLS)
- Transparent Generic Framing Procedure (GFPT)
- Asynchronous Transfer Mode (ATM)
- Synchronous Optical Networking/Synchronous Digital Hierarchy (SONET/SDH)

Because of this versatility, FC systems can have a broad application for a variety of uses that can take advantage of the benefits of each particular medium upon which the protocol resides.

## The 10G Inflection Point

While using FC on other protocols is interesting, perhaps no technology has intrigued people like the ability to use Fibre Channel over Layer 2 lossless Ethernet. In this way, Fibre Channel can leverage the raw speed and capacity of Ethernet for the deployments that are looking to run multiprotocol traffic over a ubiquitous infrastructure inside their data center.

Realistically, 10G Ethernet (10GbE) was the first technology that allowed administrators to efficiently use increasing capacity for multiprotocol traffic. It was the first time that we could:

- Have enough bandwidth to accommodate storage requirements alongside traditional Ethernet traffic
- Have lossless and lossy traffic running at the same time on the same wire
- Independently manage design requirements for both non-deterministic LAN and deterministic SAN traffic at the same time on the same wire
- Provide more efficient, dynamic allocation of bandwidth for that LAN and SAN traffic without starving each other
- Reduce or even eliminate bandwidth waste

How did this work? 10GbE provided a number of elements to achieve this.

First, 10GbE allowed us the ability to segment out traffic according to Classes of Service (CoS), within which we could independently allocate pre-deterministic and non-deterministic traffic without interference.

Second, 10GbE gave us the ability to pool the capacity and dynamically allocate bandwidth according to that CoS.

Third, consolidating traffic on higher throughput 10GbE media reduces the likelihood of underutilized links. How? Let’s take a simple example for instance. Suppose you have 8GFC links but are currently only using 4G of throughput. You have a lot of room for growth when you need it but for the most part, on a regular basis half of the bandwidth is being wasted.

Consolidating that I/O with LAN traffic and creating policies for bandwidth usage can mean that you would still have that FC throughput guaranteed, but also be able to use additional bandwidth for LAN traffic as well. Moreover, if there is bandwidth left over, bursty FC traffic could use all of the remaining additional bandwidth as well.

Because LAN and SAN traffic is not constant or static, despite what benchmark tests might have us believe, this dynamic approach to running multiple types becomes even more compelling when the bandwidth increases beyond 10G to 40G, and even 100G.

## The 40G Milestone

There is an old adage, “You can never have too much bandwidth.” If that’s true, then Data Centers are spoiled for choice. In this issue Scott Kipp has written an excellent article on how 16GFC is a shining example of how increases in wire speeds provide a great deal of additional firepower to the storage administrator’s arsenal.

In order to understand just how much throughput we’re talking about, we need to understand that it’s more complex than just the ‘apparent’ speed. Throughput is based on both the interface clocking (how fast the interface transmits) and how efficient it is (i.e., how much overhead there is).

In Table 1 you can see exactly how much the bandwidth threshold is being pushed with technologies that are either available today or just around the corner.

Speed Name	Clocking (Gbps)	Encoding (data/sent)	Data Rate (MBps)
8GFC	8.500	8b/10b	1600
10GFC	10.51875	64b/66b	2400
10G FCoE	10.3125	64b/66b	2400
16GFC	14.025	64b/66b	3200
32GFC	28.050	64b/66b	6400
40G FCoE	41.225	64b/66b	9600
100G FCoe	103.125	64b/66b	24000

Table 1: Bandwidth Threshold

The ability to increase throughput in this way has some significant consequences.



## What to Do With All That Bandwidth?

There are more ways to answer that question than there are Data Centers. Could you dedicate all that bandwidth to one protocol, whether it be Fibre Channel or something else? Absolutely. Could you segment out the bandwidth to suit your data center needs and share the bandwidth accordingly? Quite likely.

This is where the true magic of 40GbE (and higher) lies. In much the same way that SANs provided the ability for data centers to make pools of storage more efficient than silo’d disk arrays, converged networks allow storage networks to eliminate the bandwidth silos as well. The same principles apply to the networks as they did to the storage itself.

There are three key facets that are worth noting:

### Flexibility

The resiliency of the Fibre Channel protocol, exemplified by its easy transference from 10G to 40G to 100G Ethernet without the need for further modification, means that there is a contiguous forward-moving path. That is, the protocol doesn’t change as we move into faster speeds and higher throughput. The same design principles and configuration parameters remain consistent, just as you would expect from Fibre Channel.

But not only that, you have a great degree of choice in how your data centers are configured. Accidentally underplan for your throughput needs because of an unexpected application requirement? No problem. A simple reconfiguration can tweak the minimum bandwidth requirements for storage traffic.

Have space limitations, or a different cable for each different type of traffic you need? No problem. Run any type of traffic you need - for storage or LAN - using the same equipment and, often, on the same wire. Nothing beats not having to buy extra equipment when you can run any type of traffic, anytime, anywhere in your network, over the same wire.

## Growth

Data Centers are not stagnant, despite what we may see on topology diagrams or floor plan schematics. They can expand, and even sometimes they can contract. One thing they do not do, however, is remain static over time.

New servers, new ASICs, new software and hardware - all of these affect the growth patterns of the Data Center. When this happens, the network infrastructure is expected to be able to accommodate these changes. For this reason we often see administrators "expect the unexpected" by over-preparing the data center's networking capacity, just in case. No one can be expected to predict the future, and yet this is what we ask of our storage and network architects every day.

Because of this even the most carefully designed Data Center can be taken by surprise 3, 5, or more years down the road. Equipment that was not expected to live beyond its projected time frame is being called upon to work overtime to accommodate capacity requirement increases. Meanwhile, equipment that was "absolutely necessary" remains underutilized (or not used at all) because expected use cases didn't meet planned projections.

Multiprotocol, higher capacity networks solve both of these problems. No longer do they have to play "bandwidth leapfrog," where they have too much capacity on one network and not enough on the other (and never the twain shall meet!). Neither do they need to regret installing a stub network that winds up becoming a permanent fixture that must be accommodated in future growth because what was once temporary has now become 'mission critical.'

# GROWTH - BUDGET

## Budget

What happens when these needs cannot be met simply because of the bad timing of budget cycles? How often have data center teams had to hold off (or do without) because the needs of the storage network were inconveniently outside the storage budget cycle?

In a perfect world, storage administrators would be able to add capacity and equipment whenever needed, not just because of the dictates of budgetary timing. When capacity is pooled on a ubiquitous infrastructure, however, there no longer has to be a choice between whether the LAN/Ethernet capacity should trump storage capacity. Not every organization has this limitation, of course, but eliminating competition for valuable resources (not "either/or" but rather "and") not only simplifies the procurement process but also maximizes the money spent for total capacity (not to mention the warm fuzzes that are created between SAN and LAN teams!).

# ➤ FIBRE CHANNEL CONNECTIVITY

**Fibre Channel connectivity has been and continues to be as easy and straight-forward as plugging a lamp cord into a wall socket.**

By: Jay Neer, Industry Standards Manager, Molex and member of FCIA Board of Directors

Greg McSorley, Technical Business Development Manager, Amphenol GCS and member of FCIA Board of Directors

## Introduction

Fibre Channel connectivity has been and continues to be as easy and straight-forward as plugging a lamp cord into a wall socket. In the case of Fibre Channel (FC), the socket is a Small Form-factor Pluggable (SFP) receptacle. The receptacle consists of the connector which is covered by an Electro-Magnetic Interference (EMI) shield or cage as it is commonly known. The original single cage configuration was extended to ganged or side-by-side configurations and then to stacked and ganged configurations. These ganged configurations provide greater density and flexibility for configuring the FC sockets. A typical FC appliance could be configured with any of the examples shown in Figure 1. Other sizes of ganged and stacked and ganged cages are available.




Figure 1: FC Appliance Port Configurations

## Stability and Flexibility

The basic SFP receptacle has served as the primary and only mating connector system since 1 Giga-Bit/Second (Gb/s) FC was released back in '90's. The SFP "socket" is in fact, more universal than the wall socket as those of us that have travelled to other countries have discovered as there are many differing AC sockets out there. The SFP was originally designed as a socket to accept pluggable optical transceivers, electrical to optical converters for both FC and 1Gb/s Ethernet. It is still now the interface for 1, 2, 4, 8 and 16 GFC plus 1 GbE and 10 GbE Ethernet. The SFP socket is also the interconnect for Fibre Channel over Ethernet released in 2008.

To maintain this intermateable backward compatible form factor and mating interface, mechanical specifications have remained constant over time. A number of subtle internal design improvements have made it possible to achieve this backward compatibility. The contact design between the mating interface with the plug and the mating interface to the host Printed Circuit Board (PCB) have been refined to enable the interface to achieve higher data rates. The design within the host PCB interface below the connector solder tails has also been refined to accommodate higher data rates. As a result, each new generation of FC physical standard



specifies only the latest revision of the connector to assure maximum functionality. New connectors do work properly for previous revisions of the FC standard, but older versions of the connectors may not perform as well when used for newer revisions of FC.

End users needed a less costly solution for shorter links so Direct Attach Copper Cables (DAC) were developed to meet that need. The DAC's are copper assemblies made of high performance Twin Axial copper cable which is directly soldered to the PCB plug interface thereby eliminating the Electrical to Optical (E2O) conversion from each end of the cable assembly. As a result, users now have a cost effective solution for all links and lengths from 1 meter passive copper cables to 10K Single Mode (SM) optical links. The two DAC and optical module options are shown in Figure 2.




Figure 2: DAC and Optical Modules

### Configurability

The applications and their unique cable length requirements create the usage models for each of these mediums. This is facilitated by the use of an EEPROM in every SFP plug, whether it is Direct Attach Cable (DAC) or a pluggable optical module. Upon insertion or when the link is activated by the appliance the EEPROM is read. This EEPROM holds all the information about the module which provides the device specific information about what speeds the connection supports, and if it is a DAC or optical module. If it is a DAC it indicates if it is passive or active and what link lengths are supported. If it is a pluggable optical module it indicates if it supports Multi Mode (MM) or Single Mode (SM) fiber optic cables. This information and more helps the device set its output and input settings to optimize the signals for the best signal performance.

FC application configurations are many, as they are customized to meet the end user's needs.

These configurations all however consist of basic building blocks such as servers, switches and storage.

The number of building blocks for similar functional requirements may be unique due to available physical plant layouts for example. Another configuration consideration is whether the installation is a new facility or an addition to an existing one. These and other variables make the single connector interface an extremely valuable asset, as the "socket" is the same on both ends, and is the same whether the connection is between two new appliances or between a new one and an older one. The pluggable optical module plug and the copper cable plug are also the same from an interchangeability perspective.

As noted earlier, the distance between end points typically defines the type of cable selected for each installation. A representative Fibre Channel installation is shown in Figure 3. It shows DAC passive cabling for the shorter runs within racks, DAC active cabling between adjacent racks, and passive optical cabling between pluggable optical modules located in the appliances at the ends of the longer runs.




Figure 3: Generic FC Application Configuration

Looking forward, the Fibre Channel Industry Association (FCIA) Roadmap shows the 32GFC Standard completing in 2012 – and yes, it does use the SFP interface.

# ➤ HOW FC PLUGFESTS

## HELP PREPARE PRODUCTS FOR MARKET

By: Bill Martin, Engineer Consultant, Industry Standards, Emulex Corporation

### Introduction

With the rapid evolution of Fibre Channel (FC) and FCoE, providing customers with the best experience in interoperability is extremely important. FCIA sponsored plugfests provide the perfect environment to accomplish this task. The FCIA has been sponsoring two to three plugfests every year since 1996, with detailed test plans and reports generated for the participants. These plugfests have generated numerous clarifications to the standards as well as help participants debug products incorporating new technologies.

Interoperability is important not only with new technology, but also for backward compatibility with existing products that customers have in their installations. Participants in the plugfest bring their latest release FCoE and 16GFC products and code to test with other vendors to ensure interoperability and smooth integration for the end customer. Additionally, existing 16GFC products are included in the testing to ensure backward compatibility. In addition to validating interoperability of devices, cable vendors bring their optical and copper cables to ensure interoperability of new 16GFC devices over passive and active cables up to the limits specified in the standard.

In addition to providing basic interoperability testing, the plugfests provide: early visibility to new hardware implementations; configurations that would typically not be available; interaction with competitors, customers, and partners on a level playing field; ability to contribute to the future industry and technology directions; and the ability to disprove fiction and prove fact about product interoperability.

### Technologies covered by the plugfest

At the FCIA plugfests we test FCoE over 10G Ethernet and 8GFC and 16GFC native FC attachments. We observe communication between all interfaces: with 10G FCoE devices communicating through Ethernet switches with internal or external Fibre Channel Forwarders (FCFs) connected to FC fabrics with 8GFC and 16GFC devices. We validate that 8GFC and 16GFC devices negotiate to the highest supported speeds and operate within the error specifications of the FC specification.

The plugfest provides physical layer testing, which gives vendors the opportunity to validate that their device's transmitters and receivers meet the 16GFC FC specifications. It also allows for cable manufacturers to validate that their copper and optical cables meet FC specifications. Configurations are tested to demonstrate interoperability of 16GFC over maximum length cables from all participating vendors including both passive and active copper cables and all supported grades of optical cables.



**SERVERS - CNAs - FC HBAs - ETHERNET SWITCHES  
FCFs - FCoE STORAGE - CABLE VENDORS**

The plugfest brings together: Servers; CNAs; FC HBAs; Ethernet switches; FCFs; FC switches; FCoE storage; FC storage; and cable vendors. Each vendor brings their latest equipment to connect together to demonstrate interoperability in the overall storage solution.

### Configurations tested at the plugfest

Having an NDA allows the use of the information gained at the plugfest to only be used for the development of a company's products and not to be used in any form of marketing against competitors. The plugfest allows vendors to bring their latest products,

including pre-released products, to connect together with other products. The ability to test with these products allows vendors to see any potential issues that may cause interoperability issues when products are shipped to customers. Issues found at the plugfest may be as a result of improper implementation or misunderstandings of the underlying specifications. When ambiguities in the specifications are found, the issues are taken back to the standards committee, with suggested changes that clarify the standard and improve future interoperability.

The plugfest also connects together configurations that may not yet be supported in the industry, but are valid configurations according to the standards. By testing these configurations, developers get an early indication of issues that their products may experience as the industry deploys these leading-edge configurations. Configurations that the plugfest explores include configurations with Ethernet Data Center Bridging DCB switches at the edges of the infrastructure and multiple FCFs separated by DCB switches in the core of the fabric, as shown in figure 1.




Figure 1: Example configuration

This configuration allows the exploration of: how VE\_Ports work; how DCB switches forward FCoE traffic, including flow control operation; and how native FC devices may be accessed by or access FCoE devices. Deployment of this type of configuration demonstrates a growth path for customers who have a variety of equipment already installed and want to incrementally add FCoE components or 16GFC components.

By connecting configurations that are at the leading edge of the technology, participants have the ability to contribute to future industry and technology directions. Issues and opportunities that are found at the plugfest generate proposals into the standards bodies that define the protocols that are used for FC and FCoE. This process brings together companies to develop innovative solutions for customer configurations they would like to provide.

## CUTTING EDGE TECHNOLOGY

### Cooperation of competitors and partners

By operating under NDAs executed with the FCIA, competitors, customers, and partners all participate on a level playing field. The environment is one of engineers working together for the benefit of the industry as well as understanding how their individual products fit in this converged environment. There is a spirit of camaraderie at these events, with competitors and partners working together alike.

The participants work to define the test cases that are important to their product development and deployment. Vendors make engineering modifications to demonstrate the possibilities of the technology that is available, even when there may not be products yet available. This level of participation allows for the exploration of what is possible with the technology and how customers may benefit from future interoperability of competitor products.

## Dispelling fiction

The plugfest allows participants to come together and test configurations that outside observers claim do not work. One such example of this was rumors that 16GFC and 8GFC cannot interoperate. At the May 2012 FCIA plugfest, we were able to connect 8GFC and 16GFC devices together and demonstrate that these devices were able to negotiate to the highest supported speed of the devices connected. Additionally, we were able to connect 8GFC and 16GFC devices to a FC switch and demonstrate that flow control mechanisms within the switch allowed the devices to communicate with no errors.

As the industry continues to evolve and other myths are presented, the plugfests will continue to provide a forum to test these hypotheses and show how the technology is robust and allows forward migration of customers' environments.

## Conclusion

FCIA sponsored plugfests provide an environment that allows participants to prove interoperability with a variety of competitor and partner products, including backward compatibility with existing installed products. This ability helps the participants bring their products to market quicker as well as provide a better experience for our end customers.

The FCIA sponsors two plugfest each year, allowing vendors ample opportunity to come together to advance interoperability in the storage industry. These plugfests are coordinated not to conflict with other industry events. They are held at the University of New Hampshire Interoperability Lab in Durham, NH. The test plans and schedules are developed by the participants to meet their needs and the industry directions. Plugfests are open to both members and non-members of the FCIA, and details of the next plugfest can be found on the FCIA website at [www.fibrechannel.org](http://www.fibrechannel.org).



© Copyright Fibre Channel Industry Association  
Produced in the United States  
October 2012  
All Rights Reserved