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.

# FIBRE CHANNEL SOLUTIONS GUIDE


- FOLLOWING THE FCIA ROADMAP TO SUCCESS!
- POSITIONING STORAGE TECHNOLOGIES IN DATA CENTERS
- ABOUT STANDARDS FOR FIBRE CHANNEL
- FCoE I/O CONVERGENCE AND VIRTUALIZATION
- OVERVIEW OF TRILL
- ACHIEVING NETWORK CONVERGENCE: DEPLOYING FCoE INTO THE DATA CENTER
- IMPLEMENTING AND VALIDATING AN FCoE PILOT


## 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.

fibrechannel.org

# FIBRE CHANNEL

## DRIVING SPEED AND CONVERGENCE

## FOLLOWING THE FCIA ROADMAP TO SUCCESS!

*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 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 reflected in the roadmap.

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 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

**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*

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 13 years of accurate roadmaps come from the seriousness FCIA takes in this huge responsibility and obligation to the industry.

## 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.*

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 its guidelines.

# SPEED AND CONVERGENCE

*Continued...Following the FCIA to Roadmap Success*

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.

**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*

FCIA's Roadmap, version V13, 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). 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.**

# POSITIONING STORAGE TECHNOLOGIES IN DATA CENTERS

*Ahmad Zamer – Sr. Product Marketing Manager, Brocade*

## Introduction

Data networking and storage networking technologies have evolved on parallel, but separate paths. Ethernet has emerged as the technology of choice for enterprise data networks, while Fibre Channel (FC) became the dominant choice for enterprise shared Storage Area Networks (SANs). Ethernet and Fibre Channel continue to evolve on their way to higher speeds and new capabilities.

Most large organizations have invested in both technologies for data center needs. Ethernet provides the "front-end" Local Area Networks (LANs) linking users to enterprise servers while Fibre Channel provides the "back-end" SANs links between server and storage. Maintaining separate data and storage networks adds to the complexity of managing and maintaining data centers. As enterprises embrace virtualization to increase operational efficiencies, the next logical step appears to be converging server I/O traffic and links onto a common lossless Ethernet transport.

Migrating to a lossless and faster 10GbE transport opens the door for data center managers to converge LAN, FC SAN and IP storage traffic over one shared link. With that comes the need to understand where each of these storage technologies fit in future data centers. This document examines some aspects of deploying iSCSI, FCoE and FC in data centers with converged networking infrastructure.

## Data Center Storage Technologies

When discussing convergence in the data center, many talk about the various storage technologies as replacement for one another. That is a narrow viewpoint, because technologies, while they may overlap in certain aspects, always find scenarios where they fit best. FCoE, iSCSI and FC are no different for each has its own unique capabilities and can benefit data center operations.

The new Fibre Channel over Ethernet (FCoE) protocol aims at enabling the transport of FC storage traffic over a new lossless Ethernet medium, Data Center Bridging (DCB). FCoE runs over Ethernet and does not utilize TCP/IP; thus avoiding the pitfalls of relying on TCP flow control mechanisms. The most visible benefit of FCoE is in the area of server I/O consolidation. Today, servers use separate adapters to handle storage and networking traffic; Host Bus Adapters (HBA) carry storage traffic, while Network Interface Cards (NICs) handle LAN traffic. With FCoE, a single Converged Network Adapter (CNA) will be used to handle both storage and networking traffic. Using server CNAs will reduce the number of server adapters needed, which in turn reduces the number of I/O cables and the number of switch ports used. Reducing hardware resources simplifies severs I/O configurations and lowers cost of acquisition (CapEx) along with the associated operating costs (OpEx).

Before the introduction of FCoE, there was iSCSI, another encapsulation protocol that enabled transporting of SCSI block storage traffic over 1GbE TCP/IP networks. iSCSI delivered lower costs and enabled routing storage data over IP to remote data centers for backup. The price structure of iSCSI benefited from the fact that it utilized software drivers, lower priced NICs and standard Ethernet switches. Despite the lower costs, iSCSI saw deployments in small to medium enterprises with limited success in Fortune 1000 enterprises. The limitations were due, in most part, to the lossy nature of Ethernet, its higher latency, and the need for storage management tools in a networking environment. With the move to 10GbE and the addition of lossless features to Ethernet as part of the FCoE enablement effort, iSCSI will benefit and the potential for its deployment will improve. The most notable benefit will come from Priority-based Flow Control (PFC) that will help reduce the iSCSI flow control overhead associated with TCP/IP. Other areas affecting iSCSI performance, such as context switching and data copies will not benefit from the DCB features.

*Continued...Positioning Storage Technologies in the Data Center*

The migration to 10GbE is a double edged sword for it increases the cost of deploying iSCSI, while giving it a faster transport.

Fibre Channel remains the dominant storage technology in large enterprise data centers. Enterprise SANs continue to rely on the high performance and reliability of FC to run demanding database and transaction heavy applications. An increased ramp in 8Gbps FC is already underway, and an even faster 16Gbps technology is expected to be available in 2011. The higher protocol speeds will provide an easier route for data centers to deploy higher performance FC SANs without disruption. Higher FC performance will ensure that highly demanding data intensive applications continue to have a simple migration path.

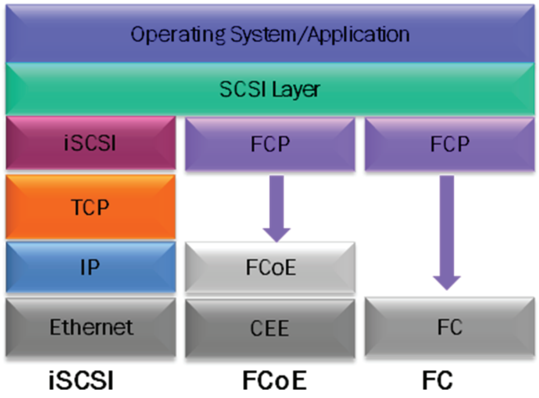


*Figure 1: Storage technologies stacks showing FCoE and iSCSI encapsulation layers as well as iSCSI use of TCP/IP*

Both iSCSI and FCoE are encapsulation protocols and are not FC replacement technologies. In fact, FCoE builds on the success of FC in the data center and utilizes FC along with new lossless Ethernet DCB, to solve servers I/O challenges facing IT professionals in data centers. Similarly, iSCSI is taking advantage of the idea of SAN that was popularized by FC success in data centers and will benefit from a wider adoption of 10GbE adoption driven by FCoE.

## Customer Segments

It is widely accepted that iSCSI found success in small and medium size enterprise, while FC is the technology of choice for large enterprise data centers. The advent of FCoE on a lossless DCB transport creates new dynamics for customers.

While lossless 10GbE benefits iSCSI by reducing TCP/IP flow control overhead, it shrinks the price advantage of iSCSI, because FCoE uses the same 10GbE components. In other words, iSCSI will run better and faster on lossless DCB, but the performance improvement will not be free and will be accompanied by a move up in prices due to 10GbE price levels. While iSCSI performance will improve benefiting from PFC, it will remain below that of native FC protocol, because 10GbE DCB does not remove other areas affecting iSCSI performance such as data copies, context switching and higher Ethernet latency. Customers should compare overall systems performances and not get distracted by results of carefully structured laboratory tests that don't reflect real life deployments or measure relevant applications performance.

It is safe to say that small businesses will continue to deploy iSCSI. On the other hand, medium sized businesses looking for new server and storage deployments will have a choice between iSCSI and FCoE. Those with existing iSCSI installed base are likely to opt for iSCSI, while ones with FC installed base will find FCoE a better fit for their needs.

Storage management will play an important role in the SAN installation decision process. It is logical to select new storage resources that utilize existing management applications, because that reduces acquisition costs and running costs associated with training and IT operations. FCoE was designed to maintain FC upper constructs allowing it to take advantage of FC management making it possible to seamlessly integrate FCoE into existing FC SAN environments. Role-based management will ease IT transition to converged server I/O links. With role-based management, networking administrators can continue to manage the LAN resources, while storage administrators continue to manage shared storage SAN resources.

Large enterprises with existing FC SANs are likely to continue adding to their FC environments taking advantage of higher speed FC links. When moving to converged infrastructures, large enterprises will lean towards FCoE for it can integrate seamlessly into existing FC installed base and can be managed using familiar applications.

In short, small businesses will continue to use iSCSI and large enterprises are most likely to continue investing in FC and FCoE, while medium size customers will have a choice between iSCSI and FCoE deployments.

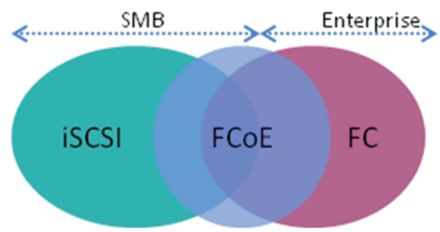


*Figure 2: FCoE, iSCSI and FC overlap, but serve different customers*

## Deployment in the Data Center

One way to appreciate where the different storage technologies may be deployed in enterprise environments is to look at three tiered data center.

As encapsulation protocols, FCoE and iSCSI will perform their functions with some overhead above that of native protocol such as Fibre Channel. As a result, data centers with genuine need for high performing FC may question the benefits of adding storage traffic to 10GbE links with LAN traffic. It is expected that FCoE and 10GbE iSCSI will most likely get deployed in environments that are currently using 4Gbps FC and 1GbE links. Customers using the lower speed I/O links will see greater benefits from migrating to 10GbE links. Benefits will come in the form of reduced I/O links and higher server utilization in virtualization environments.

iSCSI will continue to be deployed in tier three applications and may slowly migrate to a few tier two applications. It is best utilized in areas where it provides the most benefit, such as remote data backup.

FCoE is a newer technology that enterprises will initially deploy for applications where some risk can be tolerated before they expand the deployment to other areas. In the near term, FCoE will find a home in new server deployments in Windows and Linux environments with virtualized tier three and some tier two applications.

When looking at data center traffic, tier three servers providing web access generate traffic primarily made of TCP/IP data. Both 10GbE NICs and CNAs can easily service this class of servers and their related traffic. Tier two applications servers with business logic applications tend to host application of greater value to enterprises. These servers are normally connected to LANs and SANs, because their traffic is split between storage and TCP/IP. Some tier two applications, such as email and SharePoint, are good candidates for FCoE or iSCSI and would benefit from server I/O consolidation. On the other hand, tier one servers tend to host data base applications that support enterprises mission critical applications that are the backbone of business. Many of such applications service customer's needs and therefore have no tolerance for latency of disruptions. It is natural for businesses to deploy mature and reliable technologies for tier one servers and applications. In addition to that, tier one applications have a genuine need for performance and processing power that guide customers look to higher performance and reliable I/O technologies such as Fibre Channel. It is unlikely that FCoE or iSCSI will find their way into tier one environments in the near term even with 10GbE lossless transport, for speed is not the only factor in tier one environments.

| Tier 1 | Tier 2 | Tier 3 |
|---|---|---|
| Database Servers | Business Logic Applications | Web Access Servers |

**Model Tiered Data Center**

| Sample Applications | Tier 1 | Tier 2 | Tier 3 |
|---|---|---|---|
| Billing System | X | | |
| Inventory Systems | X | | |
| Research | X | X | |
| Email | | X | X |
| Test/Development | | X | X |

*Figure 3: Model tiered data center and sample applications*

## ABOUT STANDARDS FOR FIBRE CHANNEL

*Steve Wilson – INCITS TC T11 Chairman, Brocade*
*Herbert Bild – SNIA Europe German Country Committee member, NetApp*
*Gilles Chekroun – SNIA Europe Board of Directors, Cisco Systems*

In recent years, server virtualization technology started a revolution in the datacenter. Its next phase has already begun with desktop virtualization and storage virtualization. Enterprises are seeing the benefits of virtualization technology and are utilizing virtualized dynamic infrastructures.

Today's quad core machines as well as future six and eight core servers run so many virtualized machines that sometimes I/O bottlenecks develop. The obvious solution to I/O bottlenecks is to provision more I/O bandwidth by adding more adapters and switch ports and/or to migrate to higher bandwidth I/O. Current practical solutions for I/O bottlenecks include upgrades from 4GFC (4 Gigabit per second Fibre Channel) to 8GFC for storage-based services SANs as well as 1GE (one Gigabit per second Ethernet) to 10GE for Ethernet-based services LANs.

Separating Ethernet networks and FC SAN is a convenient solution for guaranteeing optimal provisioning of Fibre Channel storage-based services and Ethernet-based services. For many data center executives, maintaining separate physical networks for SANs and LANs continue to be the preferred strategy for some or all of their data center network applications now and in the foreseeable future. But with the recent advent of FCoE (Fibre Channel over Ethernet), IT executives now have the option of consolidating their tried and true Fibre Channel storage-based services AND their tried and true Ethernet-based services onto a converged DCB (Data Center Bridging) 10FCoE Ethernet network.

To address increasing user demand for higher bandwidth there have been many exciting Fibre Channel developments in the Technical Committee T11 over the past several months. These include enhancements to the FCoE protocol, and the development of the 16GFC standard. Additionally, there have been new standards developed in areas such as Simplified Configuration Management (FC-SCM), Inter-Fabric Routing, Energy Efficiency, and Switch Architecture.

This article will discuss the current status of today's published technical standards and the future developments of Fibre Channel.

### T11 - Fibre Channel Interfaces

- T11 is responsible for the standards development in the areas of Intelligent Peripheral Interface (IPI), High-Performance Parallel Interface (HIPPI) and Fibre Channel (FC).
- The standardization of IPI and HIPPI has been in progress since the mid 1980's, however, there continues to be activity on both fronts. The primary focus of T11 activities has been directed towards the Fibre Channel (FC) family of standards. It should be noted, that included in the FC family are "mappings" which allows protocols from both the IP and HIPPI standards families to be transported across fibre channel. This provides a straightforward migration path among all of the T11 standards families.
- This technical committee is the U.S. TAG to ISO/IEC JTC 1 SC25/WG4 (co-TAG) and provides recommendations on U.S. positions to the JTC 1 TAG.

incits℠
*Where IT all begins*

## Enhancements to FCoE

FC-BB-6 is currently under development in T11. FC-BB-6 describes enhancements and new functionality for the FCoE protocol. This includes support for simple point-to-point configurations and end-to-end configurations. One of the outcomes of the BB-6 development is a renewed interest in alternative topologies for standard Fibre Channel Fabrics. This will result in a new FC-SW-6 project in T11.

## 16 Gigabit per second Fibre Channel (16GFC)

The FC-PI-5 and FC-FS-3 standards define 16GFC which is the fastest edge connection and single lane interface defined for the data center. T11 is currently putting the final touches on FC-PI-5 and FC-FS-3 and it will be forwarded to INCITS for further processing and publication in August. These standards define the 16GFC architecture and describe how speed negotiation takes place between the 8B/10B and 64B/66B encoding schemes. Automatic Speed negotiation, requiring no user intervention, is the key to providing backward compatibility with the prior Base2 speeds (4GFC and 8GFC) specified by the FCIA Fibre Channel Roadmap.

## Simplified Configuration Management (SCM)

A new development in Fibre Channel is the FC-SCM effort. This project specifies the requirements for Fiber Channel support in simplified environments that are often associated with small to medium business (SMB) markets. The document focuses on behaviors associated with management tools, hosts, Fabrics, and Fibre Channel switches. The FC-SCM standard is resolving comments and will be forwarded to INCITS for further processing and publication by the end of this year.

## Inter-Fabric Routing

The Fibre Channel Inter-Fabric Routing (FC-IFR) standard defines how devices on different Fabrics may communicate without merging the Fabrics together. This simplifies the management and configuration of large configurations that may result due to the consolidation of data centers. Inter-Fabric Routing allows this communication through the use of a new Fibre Channel entity known as a Fibre Channel router. The Fibre Channel router, along with zoning enhancements, controls the manner in which devices discover one another and ultimately communicate. The FC-IFR standard is complete and has been forwarded to INCITS for further processing and publication by the end of this year.

## Fibre Channel Port Model

Due to the prevalence of virtualized environments in today's data centers, Fibre Channel is being enhanced to enable these environments. Historically Fibre Channel has provided virtualization mechanisms such as N_Port ID Virtualization, frame tagging, and Virtual Fabrics to support SAN environments. In order to unify the Fibre Channel virtualization methods, an updated Fibre Channel port model was introduced in FC-FS-3 and FC-SW-5.

Prior to the port model update, the basic link level functionality was defined as the FC-2 level. To accommodate different virtualized environments, the FC-2 level was divided into three new sublevels. The Level 2 sublevels are the Physical (FC-2P), Multiplexer (FC-2M) and Virtual (FC-2V) sublevels . The FC-2P sublevel defines low level functions such as frame transmission and reception, and buffer-to-buffer flow control. The FC-2M sublevel includes the specification of addressing and routing functions. Finally the FC-2V sublevel presents the environment necessary to support multiple high level FC-4 mappings such as SCSI and FICON. By dividing the FC-2 level into three unique levels, the FCoE environment defined in FC-BB-5 is accommodated as well. The FCoE Entity simply replaces the FC-2P and FC-2M levels while providing the proper functionality to the FC-2V level.

*Continued...About Standards for Fibre Channel*

## Future Fibre Channel Developments

T11 is continuing to develop aspects of Fibre Channel to protect customer investment and bring Fibre Channel based solutions into new markets. These advances include higher speeds, advanced protocols, energy efficiency, and FCoE refinements. Here are some specific examples:

The FC-PI-6 project was accepted in June 2010 to define 32 Gigabit per second Fibre Channel (32GFC). The FC-SW-6 project was accepted in June to define additional Fibre Channel Switch functions to support advanced Fabric topologies. These advances include support for Fibre Channel Fabric and FCoE integration, and the possible extension of FSPF for distributed FC Forwarding topologies. On the management side, FC-GS-7 project has been created to provide Fibre Channel service enhancements for FCoE and Fibre Channel Fabric environments.

T11 is also discussing the task of defining requirements for energy efficient Fibre Channel to address the green aspects of data center design and implementation. This effort will likely result in one or more additional Fibre Channel projects being created in the next few months.

In summary, Fibre Channel standards continue to evolve to meet the needs of ever-expanding SANs. Several projects are just completing including the standards for 16GFC, Inter-Fabric Routing and Simplified Configuration Management. The Fibre Channel standards have been enhanced to improve the adaptability of Fibre Channel to virtual environments and FCoE. Work has begun to enhance the basic Fabric models to accommodate new topologies in FC-SW-6 and management schemes in FC-GS-7. Work has also begun on 32GFC in FC-PI-6 to ensure that Fibre Channel remains as the fastest single-lane interface in the data center. Additionally, IEEE standards are emerging in the area of Data Center Bridging (a must for FCoE transport), and increased physical layer Ethernet switch-to-switch bandwidth for core inter-switch links (ISL) in the form of 40GFCoE and 100GFCoE ISL's.

# FCoE I/O CONVERGENCE AND VIRTUALIZATION

*Sunil Ahluwalia - Sr. Manager, Product Marketing, Intel Corp & David Berry - Sr Strategic Marketing Manager, Broadcom*
*Joy Jiang - Product Line Manager, JDSU & Ahmad Zamer – Sr. Product Marketing Manager, Brocade*

## Virtualization and FCoE

Server virtualization continues to offer data center TCO benefits while improving agility in terms of application and storage availability. Virtualization was enabled in part by the prevalence of high performance Fibre Channel SANs in Fortune 1000 data centers, and Fibre Channel has recently been enhanced to help further enable virtualization in the data center. This includes the development of the Fibre Channel over Ethernet (FCoE) protocol, which transparently leverages Fibre Channel's installed base of SAN upper layer management.

Implementing FCoE with virtualization is one way to extract value from the virtualization framework, and can improve data center cost efficiency through I/O convergence. However, the demand for storage continues to be relentless, and is growing even faster as the occupancy rate and footprint of virtual machines (VMs) increase. This is causing available bandwidth to fill up fast, and is fueling the rapid migration from 4Gbps to 8Gbps Fibre Channel and from 1Gbps Ethernet to 10Gbps Ethernet.

Most servers are hosting more and more VMs. For example, many multi-core servers have as many as 16 VMs. A February 2010 Dell'Oro Group report forecasts that the number of VMs will double between 2010 and 2012 to 16 million VMs. This spectacular VM growth will continue to drive an endless thirst for higher and higher LAN and SAN I/O bandwidth, which in turn drives the opportunity for more and more VM growth.

Consider that hypervisor vendor best practices for network I/O configurations recommended separate, redundant

Continued...FCoE I/O Convergence and Virtualization

physical 1Gbps ports for each workload within a virtualized platform. This requirement may enable the proliferation of cable and server adapter (HBAs and NICs) sprawl. However, as data centers incorporate higher I/O bandwidth, such as lossless 10GbE Data Center Bridging (DCB) technology, virtualized servers can benefit from better cost-per-gigabit economics by maximizing I/O efficiency with SAN and LAN convergence via 10Gbps FCoE.

Deployment of 10Gbps FCoE networks in converged I/O environments will utilize components that function as Fibre Channel HBAs and Ethernet NICs with a new type of server adapter called a converged network adapter (CNA) that is capable of handling both storage I/O and LAN networking traffic.

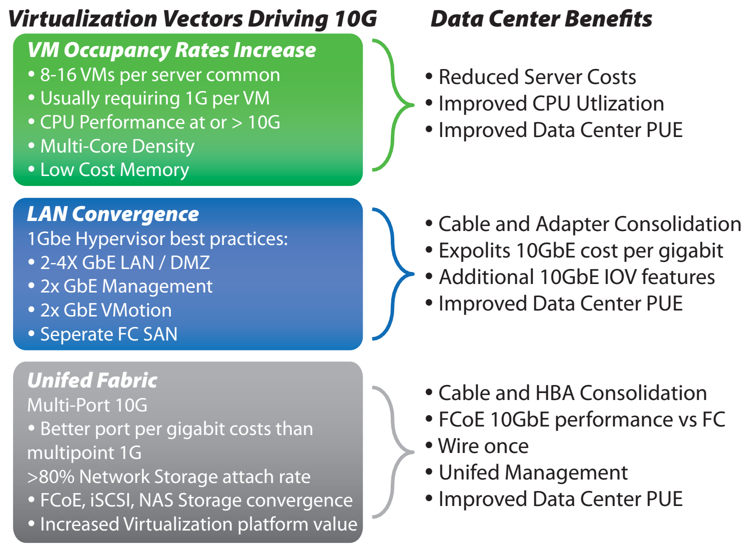


*Figure 1: Virtualization vectors driving a 10GbE unified fabric*

CNAs with lossless 10GbE provide another option in the deployment of next-generation data centers because of the potential efficiency gains from a unified fabric using FCoE-based networks.

Reliability and availability are key attributes of a SAN, and FCoE supports N_Port ID Virtualization (NPIV). FCoE zoning configurations and all virtual port properties (WWN, etc.) migrate to the new host with NPIV. A software-based abstraction layer and FCoE provide centrally managed logical pools. The virtualized FCoE network storage architecture then enables VMs to migrate across the SAN, resulting in additional availability and redundancy.

## FCoE and DCB

FCoE and lossless (DCB) Ethernet are the key technologies enabling LAN/SAN I/O convergence onto a shared I/O transport. FCoE provides seamless integration with existing Fibre Channel SANs and maintains enterprise-class availability. Data center managers can realize the maximum benefits by moving lower speed LAN and SAN traffic to the new lossless 10GbE transport.

The FCoE lossless transport needs to be immune to noise and delay. Consider that a 1E-12 BER translates to one error every 100 seconds with a 10bps data rate. Since a single bit error can cause an entire coding block (3,250 data bits) to be lost, frames could be dropped and multiple I/Os affected. Fibre Channel uses a credit-based flow control mechanism that guarantees delivery and has proven to be superior to the TCP flow control mechanism used by Ethernet. TCP is notoriously prone to data loss, and is unsuitable for transporting storage traffic. To overcome the losses inherent with TCP, lossless Ethernet uses Priority-based Flow Control (PFC) with DCB. As a result, FCoE flow control is expected to provide significant improvement over that of TCP.

As rack and row infrastructure consolidation emerges, FCoE CNAs will be deployed along with Fibre Channel HBAs and traditional Ethernet NIC cards. Layer-2 multi-path capabilities will still be able to take advantage of the installed Fibre Channel and 10GbE cabling.

## Validating FCoE and DCB

Testing and validating the Ethernet infrastructure is a key step to ensure that noise, skew and crosstalk will not be a problem. The main focus of network convergence is to access FC-based SANs through Ethernet links while

*Continued...FCoE I/O Convergence and Virtualization*

integrating host adapter functions to reduce the number of network components required for LAN, SAN and HPC applications. The role of FCoE and DCB protocols is to enable a unified fabric that effectively and reliably carries Fibre Channel SAN traffic over Ethernet. This means that in order to ensure enterprise-class performance, network operators need to take a storage-centric approach, rather than a LAN-centric approach, to test and verification of FCoE and the unified fabric.

For example, LAN testing focuses only on the switch and network. SAN testing, in contrast, requires network-to-end-device verification across all hosts, fabrics and targets contributing to the overall performance of the network. Given that the LAN fabric is based on best effort delivery, while the SAN has complete traffic control end-to-end governed by Fibre Channel's protocol for zero frame loss, LAN QoS testing is not relevant to storage networks. SAN testing measures the flow management of Fibre Channel links through mechanisms such as buffer-to-buffer credit management from the host through the fabric to storage. In addition, since SAN testing must verify the delivery of every single frame, SAN testing requires an entirely different set of performance and latency measurements compared to LAN testing.

The following should be key focal points for validation:

It is important to not confuse wire speed in lab tests with overall system performance. While pure protocol speeds are

- Protocol compliance
- Functional verification
- Performance and benchmarking tests against service level agreements (SLAs)
- Converged networks with simultaneous workloads
- Seamless integration with existing infrastructures
- Congestion management (PFC/ETS behavior)
- Unified management tools monitoring station-to-station activities

important, they only represent ideal operating conditions. The overall performance of the deployed infrastructure and applications is the real measure to take into account when building new data centers or expanding existing ones.

The critical differences between a SAN and an Ethernet LAN are the tight link-level flow control, link service management, security algorithms in use, and associated proprietary implementations. Among all of the validations, it is critical to perform interoperability tests, from the introduction of new protocols all the way through their mass deployment.

The Fibre Channel Industry Association (FCIA) continues to host plug fest events to validate interoperability between various network components. (The 4th FCIA FCoE Plugfest will take place in June at the University of New Hampshire.) Testing has been focused on protocol compliance, verifying smooth integration with Fibre Channel SANs and Ethernet, confirming lossless transport, and verifying convergence functionality. The results from these plug fests have helped to expedite interoperability between vendors to facilitate adoption of FCoE and unified fabric technologies in the end-user community.

## Conclusion

**Server virtualization, the opportunity to consolidate LAN and SAN traffic, and the increased requirement for application migration will continue to drive the need for increased I/O bandwidth within the data center. Server I/O consolidation based on FCoE and DCB lossless Ethernet is one step to achieve maximum I/O efficiency with reliable I/O consolidation. With potential reductions in capital and operating expenditures that result from efficient I/O bandwidth management, IT managers will have the option of introducing FCoE in data centers alongside existing Fibre Channel environments. The nature of FCoE as an encapsulation protocol will guide its deployment in tier 3 and some tier 2 environments, leaving Fibre Channel as the primary storage technology for tier 1 applications in the foreseeable future.**

# OVERVIEW OF TRILL

Mike McNamara – Sr. Product Marketing Manager, NetApp & Ahmad Zamer – Sr. Product Marketing Manager, Brocade

## Introduction

Many IT organizations operate purpose-built multiple networks to connect to servers. Such networks are dedicated to IP networking, shared storage, and interprocess communication (IPC) for high-performance computing environments (Figure 1). Most often overlapping networks contribute to IT organizations' added cost in numerous ways such as additional capital equipment, cable complexity, administrative costs, and additional power and cooling from multiple components.

The concept of I/O consolidation and unification or convergence enables the ability of network adapters, switches, and storage systems to use the same shared Ethernet physical infrastructure to transport different types of protocol traffic. For the IT network manager, I/O consolidation equates to installing, operating, and maintaining a single network infrastructure instead of three, while still having the ability to differentiate between traffic types. The data center manager can purchase fewer host bus adapters and NICs, cables, switches, and storage systems, reducing power, equipment, and administrative costs.

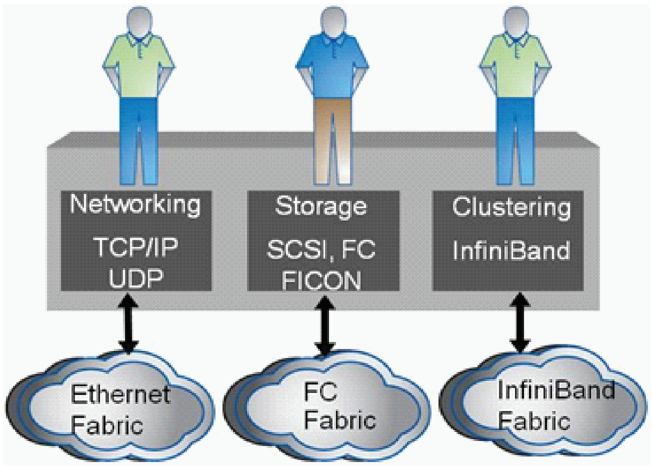


Figure 1: Purpose-bulit networks

## Lossless 10-Gigabit Ethernet

I/O consolidation and unification promise to support both storage and TCP/IP network traffic on a single network infrastructure. One of the primary enablers of I/O consolidation is lossless 10-Gigabit Ethernet, a technology with bandwidth, data integrity, and latency characteristics sufficient to support multiple traffic flows on the same link. The following factors are driving adoption and the eventual ubiquity of 10GbE as a shared I/O transport:

- Server virtualization enables application workload consolidation, which contributes to higher network throughput demands and higher bandwidth utilization.
- Virtualization aggregates multiple applications and OS instances on a single physical server with each application and OS instance generating significant I/O traffic. This places an overwhelming demand on existing multiport GbE infrastructures.
- Multisocket, multicore server technology supports higher workload levels, which demand greater throughput from IP networking and Fibre Channel storage area network (SAN) interconnections.
- Increasing use of network storage requires higher bandwidth links between servers and storage.

## Ethernet Enchancements: Standards Overview

For 10GbE to be a strong viable option for server I/O consolidation and storage networking, enhancements must be added to Ethernet to perform functions of other transports and open the door for converging multiple fabrics onto a single shared lossless Ethernet networking transport.

*Continued...Overview of TRILL*

The diverse nature of the technologies needed to enable convergence requires the development of several new industry standards that cover Fibre Channel, Ethernet, and Link Layer (Layer 2) routing. The FCoE, L2 routing, and Data Center Bridging (DCB) protocols are being developed by three different industry standards bodies, each focusing on technology areas that fall under a specific domain of expertise.

### INCITS TECHNICAL COMMITTEE T11: FCOE

FCoE (Figure 2) and FIP are defined in FC-BB-5, which describes how the Fibre Channel protocol is encapsulated, transported, and mapped over lossless Ethernet. The T11 committee completed its technical work for FC-BB-5 in June 2009 and forwarded the draft standard to INCITS, where it was approved and will be published soon.

### IEEE: DATA CENTER BRIDGING

*Figure 2: 10Gb FCoE Cable*

The Data Center Bridging effort undertaken by an IEEE 802.1 work group is aimed at adding new extensions to bridging and Ethernet so that it becomes capable of converging LAN and storage traffic on a single link. DCB is designed to make Ethernet more like Fibre Channel, because the new features being added to Ethernet are solving issues that Fibre Channel faced in the past and successfully resolved. IEEE is expected to complete its work on the various components of DCB in the second half of 2010. The new enhancements are:

#### 802.1QBB: PRIORITY-BASED FLOW CONTROL (PFC)

- Establishes eight priorities for flow control based on the priority code point field in the IEEE 802.1Q tags. This enables control over individual data flows on shared lossless links. PFC allows Fibre Channel storage traffic encapsulated in FCoE frames to receive lossless service from a link that is being shared with traditional LAN traffic, which is loss-tolerant.
- PFC provides link-level congestion control that alleviates flow control difficulties common in TCP/IP environments.

#### 802.1QAZ: ENHANCED TRANSMISSION SELECTION (ETS)

- ETS provides the capability to group each type of data flow, such as storage or networking, and assigns an identification number to each of the traffic class groups. The value of this new feature lies in the ability to manage bandwidth on the Ethernet link by allocating portions (percentages) of the available bandwidth to each of the groups. Bandwidth allocation allows traffic from the different groups to receive their target service rate (such as 8Gbps for storage and 2Gbps for LAN). Bandwidth allocation provides quality of service to applications sharing a common transport medium.
- ETS incorporates Data Center Bridging Exchange (DCBX), a discovery and initialization protocol that discovers the resources connected to the Enhanced Ethernet cloud and establishes its limits. DCBX distributes the local configuration and detects the misconfiguration of ETS and PFC between peers. It also provides the capability for configuring a remote peer with PFC, ETS, and application parameters. The application parameter is used for informing the end station which priority to use for a given application type (for example, FCoE, iSCSI). DCBX leverages the capabilities of IEEE 802.1AB Link Layer Discovery Protocol (LLDP).

#### 802.1QAU: QUANTIZED CONGESTION NOTIFICATION (QCN)

- This end-to-end congestion management mechanism enables the throttling of traffic at the edge nodes of the network in the event of traffic congestion.

**IETF: TRILL**

Internet Engineering Task Force (IETF) is developing a new shortest path frame routing protocol in multihop environments. The new protocol is called Transparent Interconnection of Lots of Links (TRILL) and is expected to be completed in late 2010:

- TRILL provides a Layer 2 multipath alternative to the single-path and network bandwidth-limiting Spanning Tree Protocol (STP) currently deployed in data center networks.
- TRILL also provides Layer 2 multihop routing capabilities that are essential for expanding the deployment of DCB/FCoE solutions beyond access layer server I/O consolidation and into larger data center networks.

## Why TRILL?

The ever increasing adoption of virtual environments in data centers necessitates a more resilient L2 networking infrastructure. Efficient and reliable L2 infrastructure is needed to support the I/O demands of virtual applications, especially when applications are migrated across servers or even different data centers. Today's STP-based networks limit the available network bandwidth and fail to maintain reliable, complex network architectures. Although there may be many Equal Cost Multiple Paths (ECMPs) or physical paths through the network at any given time, all traffic will flow along the path that has been defined by a spanning tree that includes all network devices and nodes. By restricting traffic to this tree, loops in the logical topology are prevented at the expense of blocking alternative network paths.

While STP solves the problem of traffic loops, it prevents network capacity from being fully used. Algorithms that calculate this spanning tree may take a considerable amount of time to converge following changes in the status of the configurations. During that time, the regular flow of traffic must be halted to prevent the type of network saturation described earlier. Even if multiple simultaneous spanning trees are used for separate VLANs to better distribute traffic flows, the traffic in any one VLAN will still suffer from the same disadvantage of not being able to use all of the available capacity in the network.

TRILL will enable multipathing for L2 networks and remove the restrictions placed on data center environments by STP single-path networks. Data centers with converged networks will also benefit from the multihop capabilities of TRILL Routing Bridges (Rbridges), which enable multihop FCoE solutions.

## What is TRILL?

To eliminate the restriction of STP single path through the network, the IETF formed a working group to study and solve this problem. In summary, the group was charged with developing a solution that:

- Uses shortest path routing protocols as opposed to STP
- Works at Layer 2
- Supports multihopping environments
- Works with an arbitrary topology
- Uses an existing link-state routing protocol
- Remain compatible with IEEE 802.1 Ethernet networks that use STP

*Continued...Overview of TRILL*

The result was a protocol called TRILL. Although routing is ordinarily done at Layer 3 of the ISO protocol stack, by making Layer 2 a routing layer, protocols other than IP, such as FCoE, can take advantage of this increased functionality. Multihopping allows specifying multiple paths through the network. By working in an arbitrary topology, links that otherwise would have been blocked are usable for traffic. Finally, if the network can use an existing link-state protocol, solution providers can use protocols that have already been developed, hardened, and optimized. This reduces the amount of work that must be done to deploy TRILL.

## What TRILL Does and Does Not

Although TRILL can serve as an alternative to STP, it doesn't require that STP be removed from an Ethernet infrastructure. Most networking administrators can't, and will not, rip and replace their current deployments just for the sake of implementing TRILL. So hybrid solutions that use both STP and TRILL are not only possible but also will be the norm for at least the near-term future. TRILL will also not automatically eliminate the risk of a single point of failure, especially in hybrid environments. The goals of TRILL are restricted to those listed earlier.

Another area where TRILL is not expected to play a role is the routing of traffic across L3 routers. TRILL is expected to operate within a single subnet. While the IETF draft standard document mentions the potential for tunneling data, it is unlikely that TRILL will evolve in a way that will expand its role to cover cross-L3 router traffic. Existing and well-established protocols such as Multiprotocol Label Switching (MPLS) and Virtual Private LAN Service (VPLS) cover these areas and are expected to continue to do so.

## Conclusion

TRILL is a new draft standard being created by IETF and is scheduled to be completed later this year or early 2011. The goal of TRILL is to create an L2 shortest path robust multipath and multihop routing protocol that eventually will replace the limited and restricted L3 STP.

### FC - FCoE Icons

**The FCIA Roadmap Committee assessed a need for both a Fibre Channel (FC) icon and a Fibre Channel over Ethernet (FCoE) icon to use in various production and manufacturing applications for FC and FCoE products.**

**The FCIA marketing committee worked with the Roadmap committee to develop these icons, have them trademarked, and have made them available for general industry use.**

FC Icon

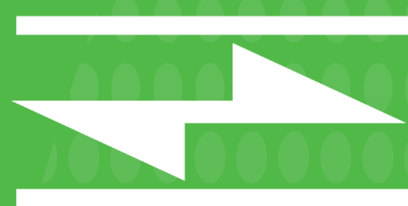

FCoE Icon

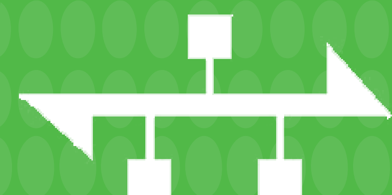

# ACHIEVING NETWORK CONVERGENCE: DEPLOYING FCoE INTO THE DATA CENTER

*Mike McNamara – Sr. Product Marketing Manager, NetApp & Ahmad Zamer – Sr. Product Marketing Manager, Brocade*

## Introduction

Many enterprise data centers rely on Ethernet for their LAN and data traffic and on Fibre Channel (FC) networks for their storage infrastructure. With increased adoption of 10GbE in the data center, the availability of Fibre Channel over Ethernet (FCoE), and new lossless 10GbE technologies, it is now possible to consolidate FC storage SAN data flows with LAN and data traffic onto the same unified Ethernet cable. Network convergence will enable enterprises to preserve their existing investments in FC storage, reduce data center costs and complexity, and simplify network management.

Although the benefits of using FCoE are quite compelling, many companies are still waiting to deploy the technology. This article provides a summary of FCoE and addresses many unanswered questions about the technology. It concludes with information on how enterprises can make the move to FCoE using a gradual, phased approach.

## IT Challenges: Maintaining Multiple Networks

The majority of today's data centers maintain multiple networks for separate purposes. Most use:

- Ethernet for their local area networks (LANs) to transfer small amounts of information across short or long distances or in clustered computing environments. Ethernet provides a cost-effective and efficient way to support a variety of data types, including corporate LANs, voice-over-IP telephony, and storage with NFS, CIFS, and iSCSI.
- Fibre Channel is used for storage area networks (SANs) to provide access to block I/O for applications like booting over SANs, mail servers, and large data-intensive databases. FC SANs are an excellent solution for storage consolidation, centralized storage management, high performance, reliability, and business continuance. The nature of FC networks as single-purpose (storage only) may result in higher costs versus general-purpose network technologies like Ethernet.

Ethernet IP data networks and FC SANs each fill an essential role in the data center, but they are very different in design and functionality. The two networks have their own security constraints and traffic patterns and utilize separate management toolsets. As a result, each network must be built and maintained on its own dedicated, isolated infrastructure, requiring separate cabling and network interfaces on each server.

Managing two separate networks for IP data and storage adds complexity and costs to the data center. Enterprises are now looking for new ways to converge their IP and SAN networks to enable the data center to run more efficiently and cost-effectively while they preserve their investments in the FC infrastructure.

## Fibre Channel over Ethernet

FCoE enables organizations to transport LAN and FC SAN storage traffic on a single, unified Ethernet cable. In this way, the converged network can support LAN and SAN data types, reducing equipment and cabling in the data center while simultaneously lowering the power and cooling load associated with that equipment. There are also fewer support points when converging to a unified network, which helps reduce the management burden.
FCoE is enabled by an enhanced 10Gb Ethernet technology commonly referred to as data center bridging (DCB)

*Continued...Achieving Network Convergence*

or Converged Enhanced Ethernet (CEE). Tunneling protocols, such as FCiP and iFCP, use IP to transmit FC traffic over long distances, but FCoE is a layer-2 encapsulation protocol that uses Ethernet physical transport to transmit FC data. Recent advances and planned enhancements in Ethernet standard, such as the ability to provide lossless fabric characteristics over a 10-Gigabit link, are what enable FCoE.

FCoE delivers significant value to organizations that are looking to consolidate server I/O, network, and storage interconnects by converging onto a single network storage technology. For data centers with large FC investments, even the simplest reduction in the amount of equipment to manage can reap benefits. And sharing the same converged network fabric—from server to switch to storage—removes the requirement of dedicated networks, significantly reducing TCO by preserving existing infrastructure investments and maintaining backward compatibility with familiar IT procedures and processes.

## FCoE Components

Some of the components needed to implement FCoE in the data center include:

- Converged network adapters (CNAs). These combine the functionality of Ethernet NICs and Fibre Channel host bus adapters (HBAs) in an FCoE environment, enabling users to reduce the number of server adapters they need to buy, cut their port count, and eliminate a healthy number of cables.
- FCoE cables. There are currently two options for FCoE cables: the optical cabling generally found in FC SANs and a new type of Twinax copper cabling. FCoE twin cables require less power and are less expensive, but because their distance is limited to fewer than 10 meters, users will likely use optical cabling to go from the top-of-rack switches to the LAN.
- FCoE switches. Users can’t go to FCoE directly from their servers to their storage arrays, so they must buy switches that support FCoE/Data Center Bridging, which connect servers to corporate LANs, SANs, or FCoE storage systems. For the early adopters, that means top-of-rack switches or end or row blades where possible.
- FCoE/DCB storage systems. These are storage systems that support FCoE and converged traffic natively. There are also storage systems that support Fibre Channel to an FCoE switch and FCoE from the switch to the host servers.

## Impact on Existing Servers, Networking, and Storage

FCoE requires minimal changes to existing IT infrastructure in enterprise data centers. It is a natural evolution of Fibre Channel network, service, and protocol technology, designed to carry data over Ethernet physical and data link layers. Using Fibre Channel’s upper layers simplifies FCoE deployment by allowing coexistence with deployed FC SANs and enabling IT to leverage enterprise-proven Fibre Channel software stacks, management tools, and trained administrators. Most importantly, IT will not need to change the enterprise’s mission-critical applications in order to benefit from the performance and potential cost benefits of FCoE.

## Organizational Issues

In traditional data center environments, the storage group owns and operates the FC SAN and the networking group owns the Ethernet LAN. Since the two groups have been historically separate, introducing FCoE into the data center may bring with it possible changes to some IT practices.

Cultural, political, organizational, and behavioral concerns in data center and provisioning paradigms can present obstacles to FCoE adoption. Some new business processes and procedures may need to be adopted to ensure that proper control mechanisms are in place for FCoE networks. Purchasing patterns may also have to be modified and the reliability of what were traditionally Ethernet networks will have to be increased.

With the convergence of FC and Ethernet on FCoE, these two traditionally separate network realms overlap. Implementing FCoE requires little if any additional IT training. FCoE leverage the existing IT expertise and skill sets of the enterprise's IP data and FC teams. Role-based management features in management applications allow the FC group to continue owning and operating the SAN, and the IP networking group to continue owning and operating the data network.

## Where to Deploy

While the benefits of using FCoE are certainly compelling, many customers are still waiting to deploy the technology end to end across the enterprise. Fortunately, FCoE convergence does not require a "rip and replace" upgrade in the data center and is not a disruptive process. Moving to FCoE can be conducted in a gradual, phased approach. Early FCoE deployments will take place mostly as part of new server deployments in Windows® and Linux® environments where virtualized tier-3 and some tier-2 applications are deployed.

Considering that FCoE is a relatively new technology, initial FCoE deployment will be best suited for access layer server I/O consolidation. As storage traffic requires the new lossless Ethernet, the 10GbE transport remains in need of Link Layer (L2) multipathing and multihop capabilities. Such features are currently under development, and should become available later in 2010. These capabilities will enable the deployment of larger FCoE networks, which will expand the reach of FCoE beyond access layer server connectivity and I/O consolidation.

Best practices for determining where to deploy FCoE in the enterprise include:

- Environments that already have a Fibre Channel skill base and Fibre Channel infrastructure
- "Green-field" deployments, where new infrastructure is being introduced to accommodate data growth
- Many enterprises will begin the transition to FCoE in their tier-3 or tier-2 applications; with the experience gained in labs or other less mission-critical environments, knowledge may then be applied to tier-2 and, in some instances, tier-1 applications
- Enterprises should start implementing FCoE on the access layer server I/O consolidation side–that step may be combined with native FCoE storage deployment; extending FCoE beyond access layer servers should wait for multipathing and multihop standards to become practical

## How to Begin

Migration to FCoE can be accomplished in a gradual, phased approach, typically starting at the edge or switch, then moving on to native FCoE storage, and eventually moving deeper into the corporate network.

The following diagram depicts typical data center architecture before network convergence with FCoE. The FC SAN (illustrated by the orange line) is a parallel network requiring network ports and cabling over and above those required for the Ethernet IP LAN (illustrated by the blue line).

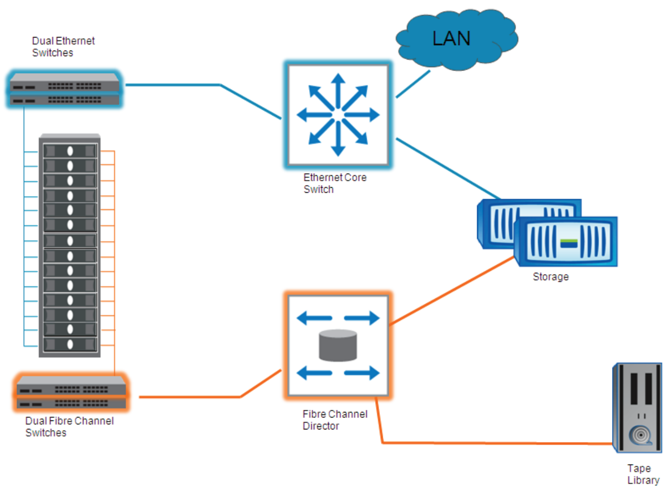


*Figure 1: Layout of a typical data center before implementing DCB/FCoE*

# SPEED AND CONVERGENCE

*Continued...Achieving Network Convergence*

## Phase 1: Transition to DCB / FCoE at the Edge or Switch

Moving to a converged or unified Ethernet infrastructure can be done gradually and will likely begin at the edge (illustrated by the green lines) where the greatest return for the investment will be realized. With FCoE convergence, port count at the servers and edge switches can be reduced by half, driving significant capital and operational cost reductions as well as improvements in management.

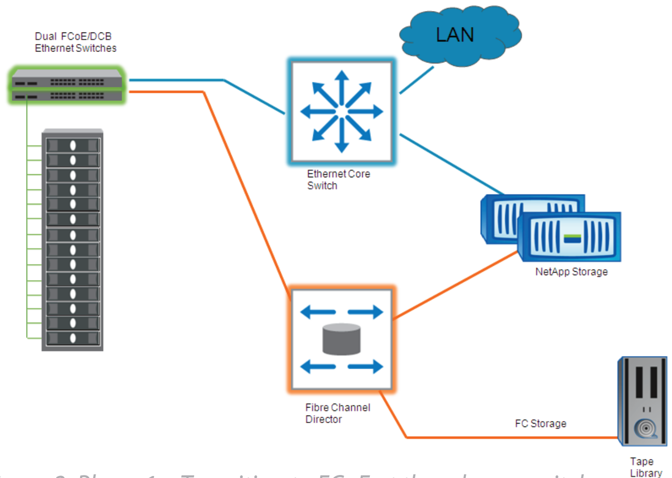


*Figure 2: Phase 1 - Transition to FCoE at the edge or switch*

## Phase 2: Transition to Native DCB / FCoE Storage Systems

Move to an end-to-end DCB/FCoE solution from the host to network to native DCB/FCoE storage. FCoE and converged traffic is supported throughout the infrastructure, providing optimal savings.

## Phase 3: Transition to DCB / FCoE at the Core

After implementing FCoE at the edge or switch, enterprises can migrate to an all 10GbE enhanced Ethernet network at the core (illustrated by the green lines) and then gradually move to supporting FCoE in the storage.

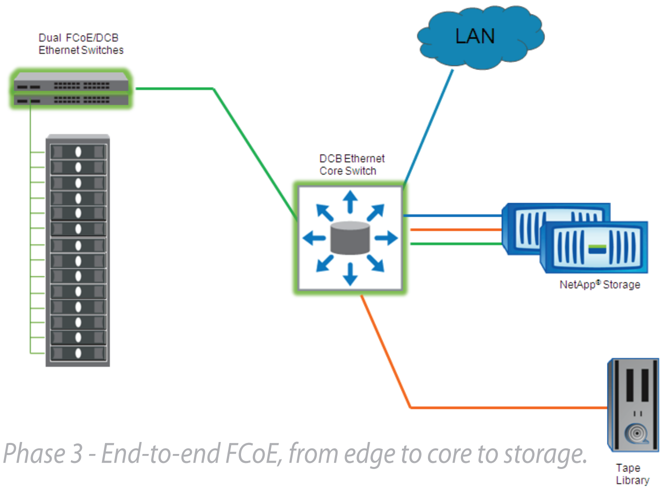


*Figure 3: Phase 3 - End-to-end FCoE, from edge to core to storage.*

## Conclusion

FCoE brings together two leading technologies—the Fibre Channel protocol and an enhanced 10-Gigabit Ethernet physical transport—to provide a compelling option for SAN connectivity and networking. And to simplify administration and protect FC SAN investments, FCoE enables IT team members to use the same tools and techniques they use today for managing both the IP and FC storage networks.

The benefits of converged data and storage networks will drive increased adoption of 10GbE in the data center. FCoE will fuel a new wave of data center consolidation as it lowers complexity, increases efficiency, improves utilization, and, ultimately, reduces power, space, and cooling requirements.

Companies that are planning new data centers or are upgrading their storage networks should take a serious look at FCoE today. By taking a phased approach to consolidating their data centers around Ethernet, customers can economically build out their Ethernet infrastructure over time while protecting previous FC infrastructure investments.

# IMPLEMENTING AND VALIDATING AN FCoE PILOT

*Joy Jiang – Product LIne Manager, JDSU & Chauncey Schwartz – Sr. Technical Marketing Manager, QLogic*

## Overview

Deploying FCoE over a unified fabric enables the convergence of data and storage networks by carrying standard TCP/IP and Fibre Channel traffic across the same high-speed 10Gbps Ethernet wire. FCoE reduces adapter, switch, cabling, power, cooling, and management requirements, resulting in substantial cost savings in the data center. A unified FCoE fabric also preserves existing Fibre Channel and Ethernet investment while providing Enhanced Ethernet for unified data networking.

## Setup

A unified fabric has two types of components:

- 10Gb Ethernet switches that implement DCB and FCoE for transporting FC traffic over 10Gb Ethernet media; also known as top of rack (TOR) switches
- 10GbE Converged Network Adapters (CNA) that replace NICs and Fibre Channel host bus adapters

A unified fabric has two types of components:

- CNA > FCoE switch > Fibre Channel switch > Fibre Channel storage
- CNA > DCB switch > FCF > Fibre Channel switch > Fibre Channel storage
- CNA > FCoE switch > FCoE storage

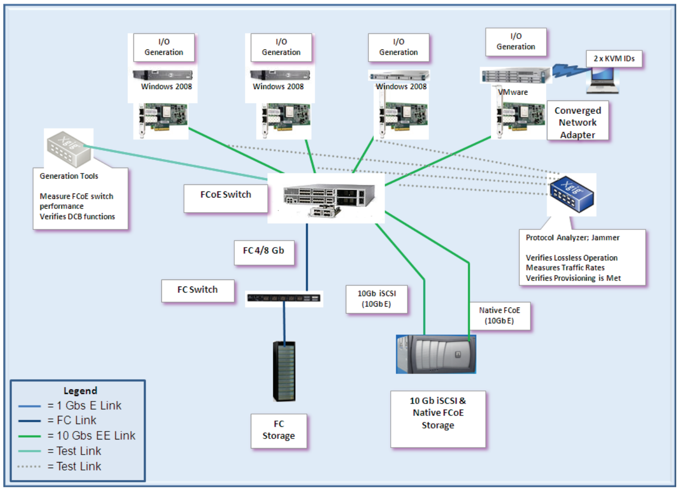


*Figure 1: Reference Architecture Diagram*

Figure 1 shows an example of a unified Ethernet network and testing infrastructure [reference X]. FCoE and iSCSI (representing the LAN) storage traffic share the unified 10GbE bandwidth driven by the CNAs. Validating the pilot network involves two major steps: Installation and verification.

## Equipment Installation and Configuration

- Install CNA Hardware: For example, install the CNA in a pilot server that meets the CNA's hardware requirements (i.e., PCI slot type, length, available slot), install the most current adapter driver and network management software, wire to the FCoE switch.
- Configure FCoE switch: Wire the pilot server to the switch port and configure the switch (i.e., port virtualization, DCB parameters, FCoE/FIP parameters if required to enable FCoE applications).
- Connect Storage: Depending upon the storage, connect directly to the FCoE switch or to a Fibre Channel switch connected to the FCoE switch.
- Verify Equipment Connections: Using network management software, verify that the server and storage

Continued...Implementing and Validating an FCoE Pilot

operate properly.

## DCB and FCoE Verification of the Unified Fabric

Ethernet is notorious for "just working", and it is essential to test and validate the infrastructure to ensure the interoperability of new technologies as well as verify the performance and efficiency of the infrastructure. Because the FCoE and data center bridging (DCB) protocols are designed for effectively carrying SAN traffic over Ethernet, verification of the FCoE/unified fabric should take a storage-centric approach (compared to a LAN-centric approach) to ensure enterprise-class deployment.

The critical differences between SAN and Ethernet LAN networks are tight link level flow control, link service management, high-security algorithms, and associated proprietary implementations. Therefore, among all the validation tests, those verifying end-to-end interoperability are the most important, from the inception of new protocols all the way through mass deployment.

### The following items are key focal points for validation:

- Protocol compliance/interoperability
- Functional verification
- Seamless integration with existing infrastructures
- Performance verification and comparison with existing non-unified infrastructures
- Impact of virtualization on infrastructure performance
- Congestion management (i.e., PFC/ETS behavior)

### To verify a unified fabric:

1) Verify PFC/ETS parameters and losslessness using application setup tests via DCBX communication
2) Validate FIP functionality, including link initialization and virtual link maintenance services
3) Verify FCoE functionality and I/O
4) Verify lossless quality of links and PFC functionality
5) Verify ETS bandwidth throttling using congestion tests
6) Measure FCoE and iSCSI converged traffic performance using interference tests
7) Measure FCoE and iSCSI converged traffic performance impact using virtualization tests

As an example, Figure 2 shows verification of a successful link initiation process between the CNA and switch FCF port and validates interoperability between the two. It also shows that the CNA has successfully established a FC-level link with FC storage across the

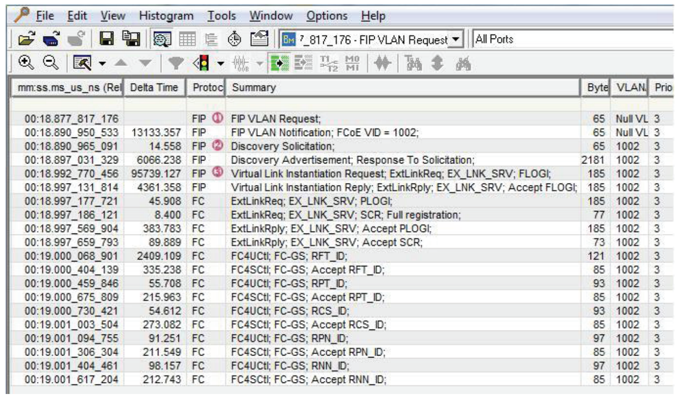


| mm:ss.ms_us_ns (Rel | Delta Time | Protoc | Summary | Byte | VLAN | Prio |
|---|---|---|---|---|---|---|
| 00:18.877_817_176 | | FIP ① | FIP VLAN Request; | 65 | Null VL | 3 |
| 00:18.890_950_533 | 13133.357 | FIP | FIP VLAN Notification; FCoE VID = 1002; | 65 | Null VL | 3 |
| 00:18.890_965_091 | 14.558 | FIP ② | Discovery Solicitation; | 65 | 1002 | 3 |
| 00:18.897_031_329 | 6066.238 | FIP | Discovery Advertisement; Response To Solicitation; | 2181 | 1002 | 3 |
| 00:18.992_770_456 | 95739.127 | FIP ③ | Virtual Link Instantiation Request; ExtLinkReq; EX_LNK_SRV; FLOGI; | 185 | 1002 | 3 |
| 00:18.997_131_814 | 4361.358 | FIP | Virtual Link Instantiation Reply; ExtLinkRply; EX_LNK_SRV; Accept FLOGI; | 185 | 1002 | 3 |
| 00:18.997_177_721 | 45.908 | FC | ExtLinkReq; EX_LNK_SRV; PLOGI; | 185 | 1002 | 3 |
| 00:18.997_186_121 | 8.400 | FC | ExtLinkReq; EX_LNK_SRV; SCR; Full registration; | 77 | 1002 | 3 |
| 00:18.997_569_904 | 383.783 | FC | ExtLinkRply; EX_LNK_SRV; Accept PLOGI; | 185 | 1002 | 3 |
| 00:18.997_659_793 | 89.889 | FC | ExtLinkRply; EX_LNK_SRV; Accept SCR; | 73 | 1002 | 3 |
| 00:19.000_068_901 | 2409.109 | FC | FC4UCtl; FC-GS; RFT_ID; | 121 | 1002 | 3 |
| 00:19.000_404_139 | 335.238 | FC | FC4SCtl; FC-GS; Accept RFT_ID; | 85 | 1002 | 3 |
| 00:19.000_459_846 | 55.708 | FC | FC4UCtl; FC-GS; RPT_ID; | 93 | 1002 | 3 |
| 00:19.000_675_809 | 215.963 | FC | FC4SCtl; FC-GS; Accept RPT_ID; | 85 | 1002 | 3 |
| 00:19.000_730_421 | 54.612 | FC | FC4UCtl; FC-GS; RCS_ID; | 93 | 1002 | 3 |
| 00:19.001_003_504 | 273.082 | FC | FC4SCtl; FC-GS; Accept RCS_ID; | 85 | 1002 | 3 |
| 00:19.001_094_755 | 91.251 | FC | FC4UCtl; FC-GS; RPN_ID; | 97 | 1002 | 3 |
| 00:19.001_306_304 | 211.549 | FC | FC4SCtl; FC-GS; Accept RPN_ID; | 85 | 1002 | 3 |
| 00:19.001_404_461 | 98.157 | FC | FC4UCtl; FC-GS; RNN_ID; | 97 | 1002 | 3 |
| 00:19.001_617_204 | 212.743 | FC | FC4SCtl; FC-GS; Accept RNN_ID; | 85 | 1002 | 3 |

*Figure 2: Three-step link instantiation process for FCoE Initialization Protocol (FIP)*

Ethernet media.

## Conclusion

Specially focusing on storage application over converged Ethernet network, the successful implementation of an FCoE pilot relies on proper interoperability of enabling protocols of FCoE and DCB and comparable performance of new network infrastructure.

## 2009/2010 FCIA Board of Directors

### FCIA Officers

- Skip Jones, QLogic Corporation: Chairman
- Tom Hammond-Doel, LSI Corporation: Vice Chairman/ Treasurer
- Greg McSorely, Amphenol: Secretary

### Members at Large

- Anthony Faustini, Cisco Systems
- Sean Fitzpatrick, Hewlett-Packard
- Rick Jooss, NetApp
- Mark Jones, Emulex
- Jay Neer, Molex
- Steven Wilson , Brocade
- Eisaku Saiki, Netmarks Inc.; FCIA Japan BOD Liaison

## 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 Solutions Guide Acknowledgements


- Alice Tate - Marketing Manager, FCIA
- Chris Lyon - Executive Director, FCIA
- David Barry - Sr. Strategic Marketing Manager, Broadcom
- Ahmad Zamer - Sr. Product Marketing Manager, Brocade
- Steve Wilson - Director, Technology and Standards, Brocade: INCITS TC T11 Chairman
- Gilles Chekroun - SNIA Europe Board of Directors, Cisco Systems
- Sunil Ahluwalia - Sr. Manager, Product Marketing, Intel Corporation
- Joy Jiang - Product Line Manager, JDSU
- Mike Tenhulzen - Alliance Marketing Manager, LSI Corporation
- Herbert Bild - SNIA Europe German Country Committee Member, NetApp
- Mike McNamara - Sr. Product Marketing Manager, NetApp: FCIA Marketing Committee Chairman
- Chauncey Schwartz - Sr. Technical Marketing Manager, QLogic
- Skip Jones - Director of Technology and Planning, QLogic: FCIA Chairman


FIBRE CHANNEL INDUSTRY ASSOCIATION



Produced in the United States

September 2010

FIBRE CHANNEL INDUSTRY ASSOCIATION

TECHNICAL AND MARKETING VOICE
OF THE FIBRE CHANNEL INDUSTRY

fibrechannel.org