BROCADE

WHITE PAPER

Understanding the Capabilities of the Emerging FCoE Protocol

Brocade is leading the standardization of FCoE, helping organizations protect and extend their Fibre Channel investments.

Table of Contents

The advent of server virtualization, the growth of digital data, and the adoption of Web 2.0 applications are driving the need for highly flexible, highly scalable, and high-performance data center architectures. And this must be accomplished while maintaining network simplicity for reduced capital and operational costs. Fibre Channel (FC) networks have long provided organizations with the reliability, performance, and network intelligence for low-latency, high-bandwidth applications.

Today, organizations are evaluating a the possibility of leveraging a new Converged Enhanced Ethernet (CEE) infrastructure, which gives them the ability to consolidate multiple transport layers on a single physical interconnect. CEE will enable a new protocol called Fibre Channel over Ethernet (FCoE), which allows organizations to extend the reliability and network services of Fibre Channel to a broader range of server environments, reducing the cost and complexity for customers implementing a virtualized data center. While FCoE is a powerful new technology that will be capable of transforming the way enterprises deploy and manage data center infrastructure, it is important to note that it is a new technology. It will take time to fully develop standards and implement correctly in a cost-efficient way. The initial cost benefits are at the server edge today and will expand to implementation of CEE and FCoE at the core of the data center switching infrastructure.

FCoE will aggregate the number of Network Interface Cards (NICs), Host Bus Adapters (HBAs), and cables needed by a server, because multiple protocols will be able to share a single physical transport layer. This will provide the framework in which to deploy a robust, powerful, scalable, and simple data center fabric. Armed with the ability to extend FC benefits to CEE networks, IT organizations will be able to support increasingly powerful applications, growing data pools, increased bandwidth, and the expansion of holistic virtualization strategies—all with managed growth and cost.

This white paper provides an overview of FCoE and CEE and the future of the new technologies. It describes the technical properties of the new protocol and transport layers and the steps IEEE, T11 and IETF standards bodies and Brocade are taking to develop these new standards. Finally it explains Brocade's evolutionary approach to adopting FCoE and building CEE networks.

Extending Fabric Intelligence throughout the Data Center

FCoE extends the transportation of the existing Fibre Channel protocol over a completely new class of transport called CEE. CEE is currently being developed by a consortium of vendors and is expected to form the basis for a series of standards in IEEE. Proponents of CEE aim to provide converged I/O for servers supporting TCP/IP, storage access, and high performance server-to-server connectivity on the same transport layer with the enhanced efficiency, security, low latency, and high performance demanded by virtual applications. The goal of this new architecture is to create a deterministic, lossless, high-performance protocol while managing growth and cost.

Today's Ethernet networks do not have these fundamental characteristics, making them unsuitable as the basis for a reliable and flexible data center fabric. On the other hand, FCoE uses the proven, standard Fibre Channel protocol and allows customers to continue to expand their existing Fibre Channel network infrastructure today and leverage FCoE and CEE technology in the future for interconnectivity—avoiding a costly and disruptive rip-and-replace migration strategy. The FCoE protocol creates a new server edge interconnect that can help reduce the cost of large-scale server deployments while maintaining the advantages of Fibre Channel.

Given the new demands put on data center architectures by the spread of server virtualization, the increase of CPU processing capabilities, additional bandwidth requirements, and the adoption of new Web 2.0 application infrastructures, organizations are looking to consolidate multiple protocols on a single transport layer to create powerful, robust, flexible, cost-efficient, and simple data center infrastructures.

Today, some rack-mount and blade servers implement a complex set of interface cards to connect to TCP/IP, InfiniBand, and Fibre Channel networks. Many organizations deploy multiple interface cards of the same protocol in each server, providing performance, redundancy, or network separation. However, this leads to servers with four, eight, or even 12 different interface cards and their corresponding cables connecting the server to the various networks. Cabling complexity becomes an issue, since each of the interface cards has at least one cable connecting it to a switch at the top of the rack or at the end of a row. CEE would allow a single interface card in the server to run multiple protocols over the same physical interface.

Client/server communication using TCP/IP, server-to-server communications using Remote Direct Memory Access (RDMA), and server-to-storage communication using Fibre Channel would all use the same physical connection but be isolated in separate Virtual Channels on CEE. This allows customers to consolidate the number of connectivity devices needed in the server to two redundant interface cards.

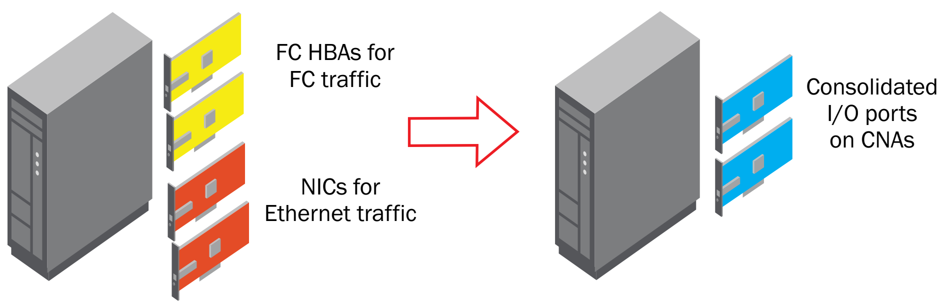

Figure 1: CNA cards replace HBAs for FC traffic and NICs for Ethernet traffic

Benefits of Utilizing CEE

The ability to use CEE as the foundation of the data center architecture provides many cost savings, including initial capital costs and on-going operational costs. These include:

- **Equipment procurement** – Organizations would be required to purchase fewer interface cards.
- **Cables and cabling cost** – Organizations would be able to consolidate the number of interfaces in the server, resulting in fewer cables from the server to embedded or top-of-rack switches. This reduces cable expenditures and management complexity, mitigating the potential for human error.
- **Power and cooling** – The consolidated interface cards require less power and cooling.

Current State of CEE Technology

However, despite the potential for cost, complexity and performance benefits through the implementation of FCoE over CEE networks, the technology is several steps away from real-life, wide-scale deployment. FCoE is still new and its foundation is still under development, and like all new technologies, the protocol needs to be tested, fully developed in standards bodies, and rolled out in phased deployments in an effort to save customers time, money, and network complexity while mitigating risk.

While the price point of the new technology is still being investigated, it is expected and widely forecasted that 10 GbE CEE switch port, optics, and server interface adapters will come down in cost and will become affordable enough that consolidation of interfaces cards and cables will make economic sense.

Another obstacle to quick deployment is that customers have traditionally deployed two separate networks for security, reliability, and performance reasons: one for connecting servers to clients and other servers (front end) and another for connecting servers to storage (back end). There is still significant development work to be done to ensure that the traffic from the front-end network does not adversely affect traffic from the back-end network and vice versa.

In addition, before FCoE is widely deployed, organizations need to understand how protocol consolidation on CEE networks would affect availability and performance of their business applications, especially in virtual environments where the physical location of an application and its associated data may not be the same ten minutes from now. How would a security attack on the TCP/IP network impact the security of the storage environment? Would the increased traffic flow of routine backups significantly degrade application performance?

Cultural, political, organizational, and behavioral concerns in data center and provisioning paradigms will also be a big obstacle to FCoE adoption. New business processes and procedures will need to be adopted to ensure that the same change control mechanisms are in place for CEE and Fibre Channel networks. Purchasing patterns may have to be altered and the reliability of what were traditionally Ethernet networks will have to be increased.

Organizations should evaluate CEE carefully, recognizing the difference between tangible benefits today and the potential benefits dependent on future advances.

Here are some examples of how the benefits of FCoE may not be realized in initial deployments and why organizations should do due diligence before deploying the new technology:

- Cables and cabling cost. Effective data rate of current pre-standard CEE interface cards is much lower than 10 Gbit/sec, requiring that organizations deploy multiple interface cards which would eliminate any cabling advantage.
- Power and cooling. Current pre-standard 10 Gbit/sec CEE interface cards often have many components to terminate the varying protocols that run over CEE, increasing power consumption.

That being said, the price of 10 Gbit/sec CEE interfaces will eventually drop, and the potential power and cooling benefits will be realized as more functionality is integrated into common ASICs on the interface cards. CEE is a promising technology, but, as with all new technologies, IT departments should carefully assess the benefits versus the inherit risks of deploying new, unproven technology. As the technology matures, FCoE will make sense as an interface to consolidate I/O for those servers with high numbers of network interfaces.

Given these inhibitors, many organizations are planning a phased migration to FCoE, starting with the server edge and then extending to the data center core. This intelligent, phased approach will allow them to try out the new technology on the edge where FCoE provides the most benefits today while still enjoying the benefits of the safer and more established Fibre Channel protocol at the more critical core. Once FCoE is fully developed, tested, and has been proven to be robust, organizations may then think about a wider implementation—but it is important to hammer out these details now.

FCoE Migration Strategy Relies on Existing Fibre Channel Infrastructure, Built-In Fabric Intelligence

FCoE is not a replacement for Fibre Channel—which will continue its vital role in the data center due to its built-in performance, scalability, and security. Gartner predicts that "for non-critical and mission-critical applications, converged networks must be engineered to the appropriate level of reliability. In life-critical applications separate networks probably will be required."

Instead, FCoE augments existing Fibre Channel fabrics, providing enterprises with additional server connectivity options. Different servers will have different requirements, which will drive the choice of protocol, fabric interface, and the overall design of the fabric itself. However, until FCoE is ready for full deployment in production environments, organizations will rely on Fibre Channel in the data center.

Until that time, careful work is underway to ensure that Fibre Channel seamlessly interoperates with the server environment whether it runs over a Fibre Channel physical interface or over a CEE physical interface. This standardization and consistency will enable organizations to eventually migrate current procedures, policies, and best practices to CEE if they so choose. For example:

- An FCoE port will be deployed and managed the same way as today's Fibre Channel ports. Administrators could simply zone the FCoE fabric the same way as a Fibre Channel fabric, using worldwide node names for addressing and standard Fibre Channel security policies and configurations.
- Organizations can utilize existing Fibre Channel management products to manage an FCoE fabric or a combined FCoE and Fibre Channel fabric.
- FCoE extends the value of SAN features and fabric services—including replication, data-at-rest encryption, and storage virtualization—across to servers with FCoE/CEE interfaces, allowing customers to leverage existing fabric intelligence and investments.
- FCoE and CEE are nascent technologies and may not be widely deployed for a number of years, allowing organizations to keep the Fibre Channel protocol layer in use today while investing in FCoE and CEE tomorrow.

FCoE and CEE: A technical perspective

The primary advantage of running the FCoE protocol over a CEE network is that it allows consolidation of I/O connectivity at the server edge, while alleviating any change to back-end storage systems, current storage management interfaces, and best practices. This key benefit provides organizations with the framework in which to build a dynamic fabric while preserving existing investments in equipment and management.

However, how is CEE networks engineered? First, it is important to understand why networks were diverged to begin with.

A fundamental assumption behind all TCP/IP-based applications is that networks are inherently unpredictable, and therefore, it is the responsibility of the endpoint, or host, to provide the required availability and reliability. This assumption has guided the design and implementation of TCP. Ensuring reliable communications between network end points has become an integral part of networking protocol stacks in host operating systems. Once TCP was widely available in the endpoints, there was never a need to engineer reliability directly into the protocol.

It is important to contrast this basic assumption of TCP/IP applications to that of high-performance applications, such as SCSI block storage and clustering applications, and what they expect out of the transport infrastructure. These applications are built assuming a reliable and highly available physical infrastructure, relying on Fibre Channel or InfiniBand to provide a network interconnect that is inherently reliable, deterministic, lossless, and high performing. This reliance on a robust protocol allows applications to focus on delivering or manipulating data, rather than on implementing a robust transport layer. Some applications developed for high-performance interconnects actually eliminate the transport layer in their stacks in an attempt to increase application performance. The proven transparency of Fibre Channel and the high-performance and reliability of block storage applications allows the responsibility for reliable transport of data to reside in the infrastructure.

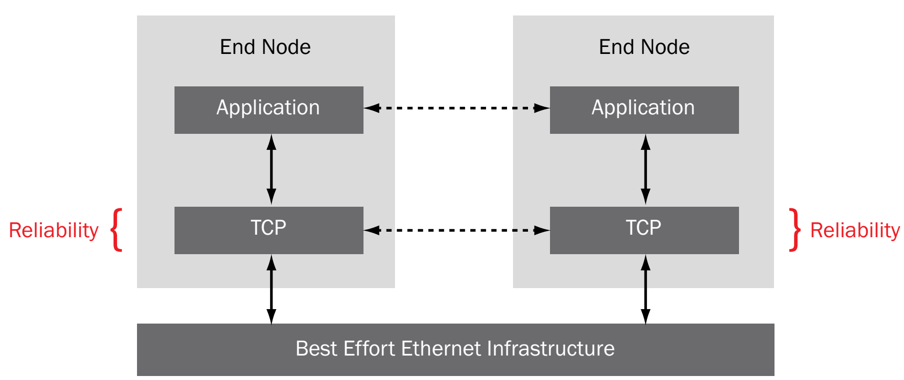

Figure 1: Reliability in end points for TCP/IP applications which wastes processing cycles and hinders application performance.

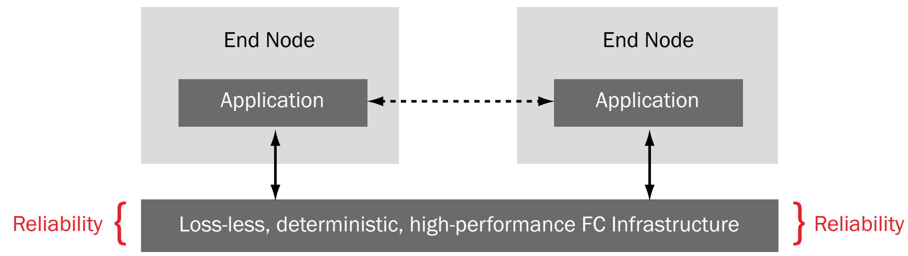

Figure 2: Common reliability in Fibre Channel infrastructure improves application performance.

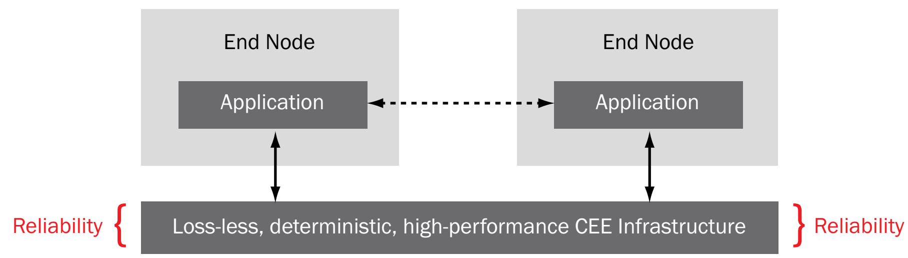

Figure 3. Common reliability in CEE infrastructure simplifies end-points.

The development of FCoE challenges the classic TCP/IP network assumption and extends the reliability and predictability of Fibre Channel to CEE networks. The vision behind CEE is to greatly enhance the IEEE-defined Ethernet link protocols with reliability and predictability functions while being backward compatible with existing Ethernet infrastructures. With reliability and performance built into the infrastructure, the goal is then to consolidate storage (Fibre Channel) and clustering (Inter-Process Communication) I/O directly on top of CEE while continuing to support typical TCP/IP applications.

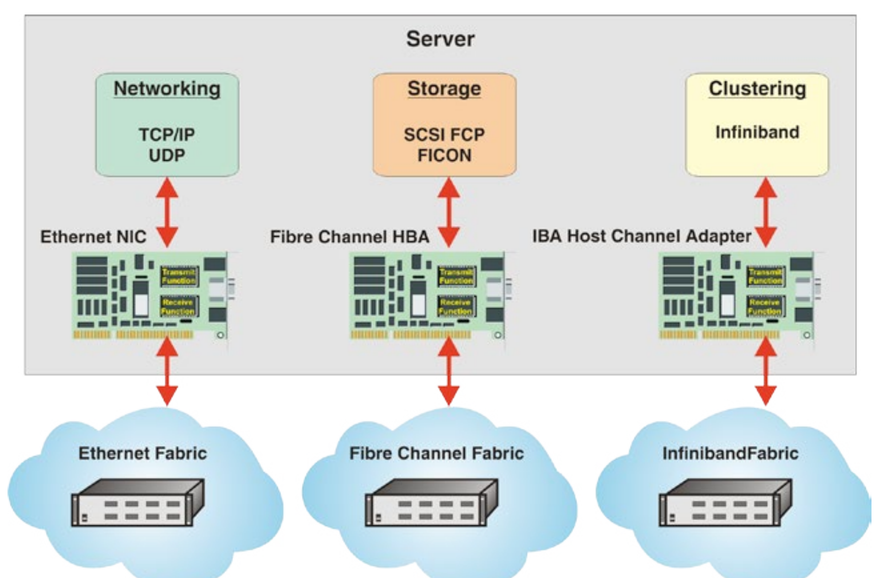

Figure 4: A traditional deployment model requires three separate network interface cards (NICs).

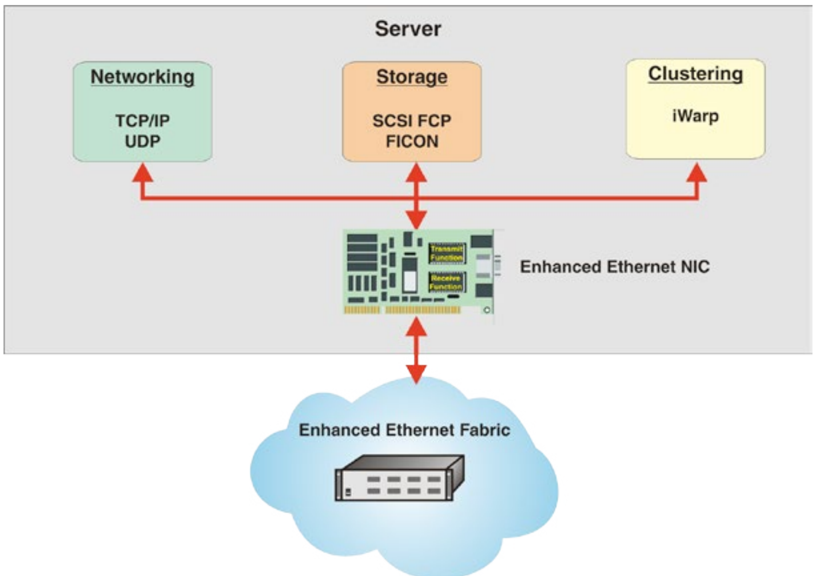

Figure 5: A CEE network allows multiple protocols to run through a single Converged Network Adapter (CNA).

Multiple enhancements are currently being considered for CEE in different working groups in the IEEE standards body with the singular goal of enabling greater amounts of reliability in networks deployed inside the physical confines of a data center. These enhancements are not required outside data centers, since all clients throughout the organization would still be relying on existing TCP/IP networks to connect to application, Web, and database servers consolidated in the data center. CEE provides the most benefits inside the data center specifically for those applications that require a reliable, robust infrastructure.

FCoE Framing

FCoE enables the transport of all existing higher-layer Fibre Channel storage protocols over CEE. FCoE encapsulates a fully-formed Fibre Channel frame and inserts a defined frame header before it is sent on the CEE link. FCoE enables a Fibre Channel operation over CEE infrastructure in a manner identical to how IP protocols currently operate.

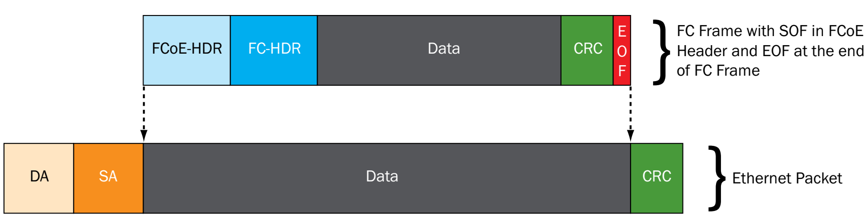

Figure 6: Fibre Channel framing inside Ethernet.

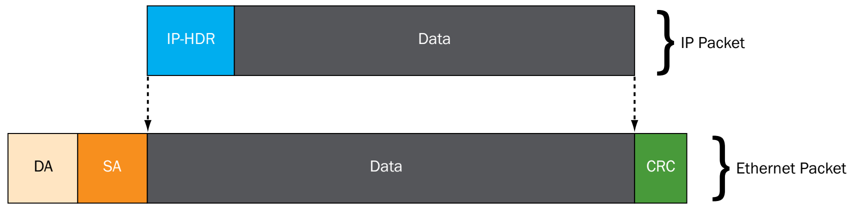

Figure 7: IP framing inside Ethernet.

FCoE encodes the Start of Frame (SOF) and includes it in the FCoE header prepended to the standard Fibre Channel frame. It then encodes End of Frame (EOF) at the end of the Fibre Channel frame before the frame is handed over to the CEE layer. A header is added and a Cyclic Redundancy Check (CRC) is completed before transmitting the unmodified Fibre Channel frame as data.

As shown in Figure 6 and Figure 7, an FCoE initiator can also be used as a Fibre Channel initiator, and the only difference is the existence of the additional link layer under the Fibre Channel protocols.

FCoE Switch

An FCoE switch is a device that exposes CEE and Fibre Channel ports, making forwarding decisions based on the Fibre Channel headers inside the CEE frame. It is similar to how a basic Fibre Channel switch would switch Fibre Channel frames. The only difference is that an FCoE switch supports communication between any ports on the switch—Fibre Channel or CEE.

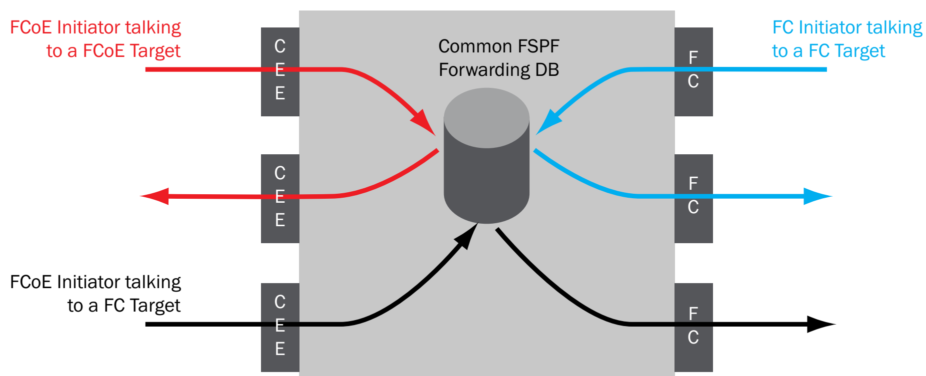

Figure 8: FCoE switch supports Fibre Channel traffic between CEE ports and between Fibre Channel ports and CEE ports.

An FCoE switch is also capable of receiving and processing basic TCP/IP traffic on the CEE ports since it is connected to a host—capable of acting as a standard IEEE 802.1Q bridge between all CEE ports in this Layer 2 mode. This operation is very similar to how a multi-layer Internet Protocol (IP) switch performs basic Ethernet switching at Layer 2 as well as IP routing at Layer 3 between every router port. An FCoE switch treats Fibre Channel traffic as a Layer 3 network protocol, and no LAN traffic is forwarded to Fibre Channel ports in the L2 mode.

An FCoE switch can be connected to other FCoE switches via the CEE ports, and it also can be connected to any existing Fibre Channel switch via the Fibre Channel ports.

FCoE in the Brocade Data Center Architecture

The Brocade data center architecture helps customers transform data centers into a flexible and dynamic environment, taking advantage of mulit-vendor, multi-platform, and multi-protocol environments. This new data center infrastructure facilitates resource optimization while protecting existing equipment investments and ensuring maximum application availability and performance.

In order to achieve these goals the Brocade DCF architecture integrates existing and future technologies into an integrated framework, which further enhances the performance, scalability, and security of the data center infrastructure. The advantages of extending the reliability, performance, scalability, and network intelligence of Fibre Channel networks to CEE has promising consolidation benefits for enterprise servers and continuing end-to-end in the data center.

With more than a decade of experience building data center architectures, Brocade is in a unique position to help organizations investigate, plan, deploy, and operate high-performance, low-latency, and highly scalable data center connectivity solutions. The Brocade strategy to enable server I/O consolidation is to utilize an evolutionary model—instead of a revolutionary one—minimizing risk and allowing storage network administrators to plan for FCoE while taking advantage of existing Fibre Channel infrastructure. In addition, FCoE over CEE networks is an important technology in enabling Brocade virtual data center environments.

The Brocade Solution: Server I/O Convergence with FCoE

Brocade offers a choice of products with the most flexibility for deploying FCoE in a current infrastructure. The products offer the ability to converge server I/O for a wide variety of servers ranging from blade, rack, and standalone systems. Brocade's initial FCoE product offerings can be deployed in the following configurations:

- Rack server I/O convergence with a top-of-rack FCoE switch
- Blade Server I/O convergence with a mezzanine adapter
- Rack server with a stand-up adapter
- Pedestal server with a standup adapter
- Blade server I/O convergence with a blade server FCoE Switch Module
- High-end server I/O convergence with an FCoE Switch Blade for the Brocade DCX™ Backbone

The FCoE Switch Module for blade servers provides CEE switching between the servers and forwards all LAN traffic to the external switch in the access layer. All FCoE traffic from the servers inside the blade enclosure is translated into Fibre Channel within the FCoE Switch Module and switched to existing Fibre Channel SAN fabrics. TCP/IP traffic is handled the same way as over Ethernet networks.

Brocade Working with IEEE to Develop FCoE Standards

IEEE is working with multiple groups to enhance the reliability of standard Ethernet so that TCP/IP applications can share the data center infrastructure with Fibre Channel and other high-performance computing applications. Each of these groups is focused on the following specific individual enhancements to the existing IEEE 802.1Q standard:

- IEEE 802.1Qbb – Priority-based Flow Control (PFC)
- IEEE 802.1Qaz – Enhanced Transmission Selection (ETS)
- IEEE 802.1 DCBX – Data Center Bridging Exchange Protocol

ADDITIONAL FCOE TERMINOLOGY

FCoE VN Port:

FCoE equivalent of a Fibre Channel N_Port

FCoE VF Port:

FCoE equivalent of a Fibre Channel F_Port

FCoE VE Port:

FCoE equivalent of a Fibre Channel E_Port

One additional enhancement feature, congestion management, will not be part of the DCB initial effort, but may be added later.

- IEEE 802.1Qau – Ethernet Congestion Management (ECM)

Priority-based Flow Control (PFC)

PFC attempts to emulate the Virtual Channel technology widely deployed in current Brocade Fibre Channel SANs. It is an enhancement to the current link-level flow control mechanism defined in IEEE 802.3X (PAUSE). Current Ethernet protocols support the capability to assign different priorities to different applications, but the existing standard PAUSE mechanism ignores the priority information in the packet. This results in a shutdown of the entire link for all applications—even if a single application causes the congestion. For FCoE, LAN, and IPC applications to effectively share the same link, it is important to ensure that congestion caused by any one of these applications does not cause disruption to the rest of the application traffic.

IEEE 802.1Qbb is tasked with enhancing the existing PAUSE protocol to include the priority in the packets contributing to congestion. Based on the priority information collected through PAUSE, the server stops sending any traffic for that specific application while the other applications continue to make progress without disruption on the shared link.

Enhanced Transmission Selection (ETS)

ETS attempts to manage the traffic priorities between multiple applications by regulating flow and by assigning preset amounts of link bandwidth and relative priority to each application. IEEE 802.1Qbb is tasked with defining scheduling mechanisms for managing traffic priorities on a converged link.

Applications can be assigned multiple priority levels—out of a total of eight—depending on the importance of packet delivery. For example, FCoE can use two different priorities–one for Class F frames and the other for regular Class 3 frames. Each application is assigned priority groups and a percentage of the maximum link bandwidth they are allowed to use. An example specification is shown below.

Priority Group ID	Bandwidth %	Application
1	-	IPC (Clustering)
2	50%	LAN (TCP/IP)
3	50%	SAN (FCoE)

Table 1: Priority Group ID numbers assigned to applications to intelligently allocate bandwidth.

When a CEE link is configured as shown above, all IPC traffic is transmitted first, and then the remaining bandwidth is equally shared between LAN and FCoE traffic. As a result, the IPC applications—which may be sensitive to latency—are never starved by the other applications.

DCB Exchange Protocol

The primary goal of DCB Exchange Protocol (DCBX) is to allow the mutual discovery of CEE capable hosts and switches and then allow CEE specific parameters—such as those for ETS and other capabilities—to be sent before the link is shared. It can also be used to define the limits of the CEE cloud and limit the applicability of Ethernet enhancements within that CEE cloud. IEEE 802.1 is tasked with enhancing existing standard Link Level Discovery Protocol (LLDP) with CEE specific parameters.

The success of FCoE adoption hinges on robust standards and interoperable solutions from an ecosystem of vendors. Brocade and other key technology partners formed CEE Authors in December 2007 to finalize the technical details around PFC, ETS, and DCBX mechanisms. The CEE Authors group reached consensus in March 2008 and submitted a joint proposal to IEEE as a “version 0” standard. Finalization of these standards is expected to be complete in late 2009 or early 2010 by a wider group of companies participating in the respective working groups. It is expected that the products implementing “version 0” of enhancements will be ready for technology trials and further refinement of the eventual standards.

Summary

Converged Enhanced Ethernet (CEE) and Fibre Channel over Ethernet (FCoE) are powerful new technologies that organizations can leverage to extend Fibre Channel reliability, network services, intelligence, and performance to networks that have traditionally relied on Ethernet. As critical components of the Brocade data center architecture, these new protocols will allow organizations to build a reliable, powerful, and dynamic infrastructure that enables data center virtualization solutions.

Once it is fully developed and robust standards have been established, CEE has the promise of reducing complexity and costs at the server edge by reducing the number of adapter cards and cables needed to connect servers to Fibre Channel, TCP/IP, and high-performance computing networks. FCoE will have the added benefit of allowing seamless operation of the Fibre Channel protocol over new CEE networks, allowing for consistent management and operation of the storage network.

FCoE and CEE are attractive new technologies for the consolidation benefits they bring. Brocade and its partners are driving the development of the CEE and FCoE standards through involvement in IEEE, positioning the company as a leader in this new, innovative technology. However, the CEE and FCoE standards are still being developed, which mandates a pragmatic approach to implementing new solutions, starting at the server edge and then extending to the data center core. Brocade recommends an evolutionary strategy, working with customers to leverage backward and forward interoperability of its server connectivity solutions until the standards are ready for full deployment. In this way, organizations can take advantage of existing Fibre Channel infrastructure, while laying the foundation for future FCoE deployments. To learn more, visit www.brocade.com.

Corporate Headquarters
San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters
Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters
Singapore
T: +65-6538-4700
apac-info@brocade.com

BROCADE