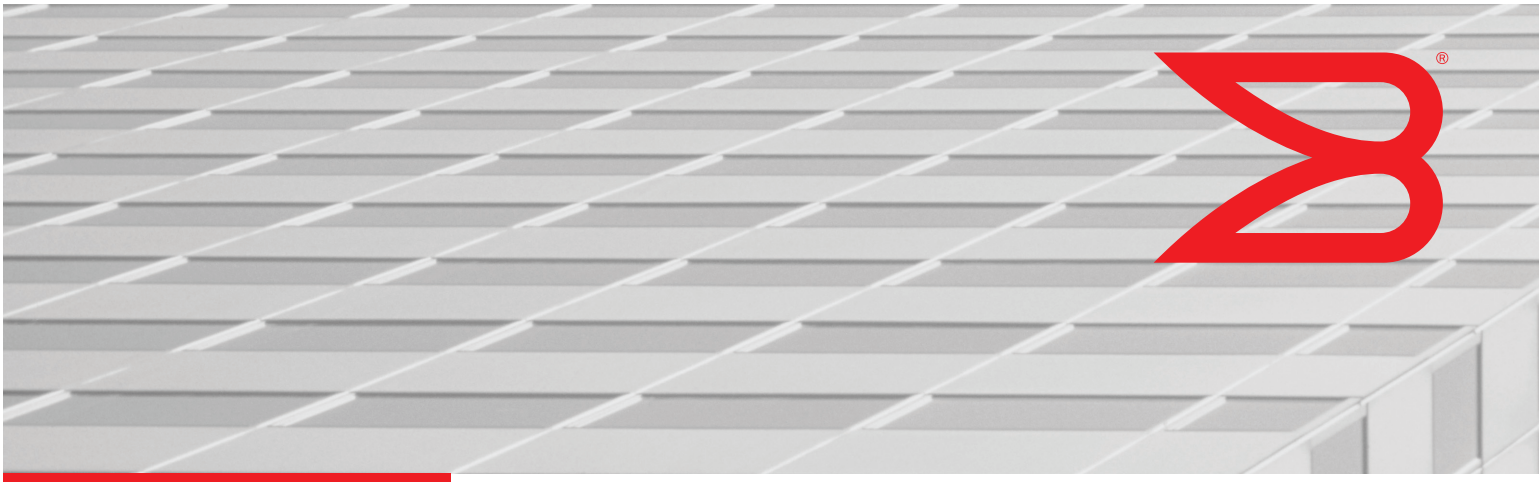

DATA CENTER

Fiber Channel over Ethernet: Enabling Server I/O Consolidation

Brocade is delivering industry-leading FCoE solutions for the data center with CNAs, top-of-rack switches, and end-of-row FCoE blades that make server I/O consolidation a reality.

BROCADE

The development of Fibre Channel over Ethernet (FCoE) opens the door to the real possibility of transporting more than one I/O traffic flow over a Converged Enhanced Ethernet (CEE), also called a Data Center Bridging (DCB), link. Moving Storage Area Network (SAN) and Local Area Network (LAN) traffic onto a single transport benefits data centers in many ways. It reduces the initial capital outlays for equipment (CapEx) and lowers the running costs of maintaining and cooling data centers (OpEx).

CEE and FCoE combine the traditional Ethernet and storage ports into a single 10 GbE CEE ports, thereby reducing the number of ports in servers. That in turn reduces cabling configurations and simplifies data center environments. The real advantage is that FCoE achieves this without disrupting existing Fibre Channel (FC) installations, as it preserves FC constructs and FC management tools. FCoE and CEE deployment will enhance the evolution of data centers by promoting more simplified cabling, more integrated management, and greener environments.

This paper provides an overview of FCoE and CEE and explains their benefits in the area of server I/O consolidation.

OVERVIEW OF CEE AND FCoE

FCoE is an encapsulation protocol that wraps FC frames into Ethernet frames and prepares them for transport over a Converged Enhanced Ethernet link. The FCoE protocol stack (see Figure 1) is constructed by taking FC upper services (Layers 2, 3, and 4) and placing them on top of Ethernet physical and Data Link layers (Ethernet). Sandwiched between the FC and Ethernet layers is the FCoE layer, which serves as a translator between the two FC and Ethernet protocols. The FCoE layer encapsulates FC to Converged Enhanced Ethernet traffic and performs the reverse function on CEE-to-FC traffic.

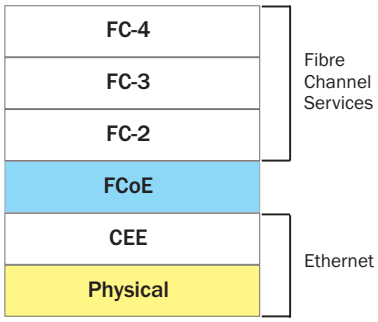

Figure 1.
FCoE stack.

As defined by industry standards bodies, FCoE simply wraps the entire FC frame as is without any modifications. The fact that the FC payload remains intact throughout its FCoE journey allows FCoE to preserve the FC constructs and services and utilize existing management applications. As a result, FCoE solutions are expected to integrate seamlessly into existing FC environments without introducing disruptions or incompatibilities. Note, however, that FCoE will not alleviate existing incompatibilities in existing FC products or environments.

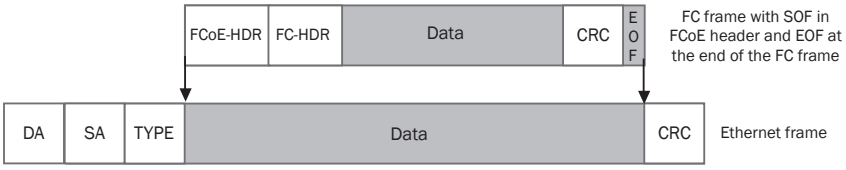

Figure 2.
Ethernet frame with FC frame insertion.

As an encapsulation protocol, FCoE builds on the success of Fibre Channel in the data center and serves to extend its presence. Contrary to some beliefs, FCoE does not compete with FC for encapsulation protocols; it supplements rather than compete with storage interface or networking protocols such as FC. It is hard to talk about FCoE without covering CEE, as each one complements the other and they need each other to deliver their promised benefits.

When distilled to its most basic form, CEE can be viewed as a set of enhancements to traditional Ethernet. Industry standards bodies have created three enhancements that provide the lossless transport of FCoE data over a CEE link. They are:

- **Priority Flow Control (PFC).** Allows the flow of higher-priority traffic to continue while lower-priority traffic may be temporarily paused when congestion occurs.
- **Enhanced Transmission Selection (ETS).** It is easier to understand this feature in terms of bandwidth management. ETS enables the grouping of each type of I/O flow on the wire into separate groups such as storage, networking, and so on. Each group is then assigned a percentage of the available bandwidth. For example, the storage group may be assigned 60 percent of the bandwidth, while LAN traffic is given 30 percent. The purpose of ETS is to provide Quality of Service (QoS) capabilities to applications that may need it and to offer bandwidth flexibility.
- **Data Center Bridging eXchange (DCBX).** A discovery and configuration protocol that allows the discovery of FC nodes connected to the CEE fabric, configuration of the nodes, and assignment of a virtual MAC (Media Access Controller) address to each note. DCBX negotiates the capabilities of the nodes during the discovery process.

Industry standards bodies will continue to add new enhancements to CEE in the near future. Two such features that are in process are congestion mangement that adds QoS to the protocol and Transparent Interconnection of Lots of Links (TRILL), which adds Layer 2 (L2) multipathing capabilities to CEE.

NEW ADAPTERS AND SWITCHES

Translating new technologies into reality always generates new products and FCoE and CEE are no exception. Traditional Host Bus Adapters (HBAs) used in storage networks and Network Interface Cards (NICs) used in Ethernet networks will be updated to a single new device, called a Converged Network Adapter (CNA).

Figure 3.
CNA architecture.

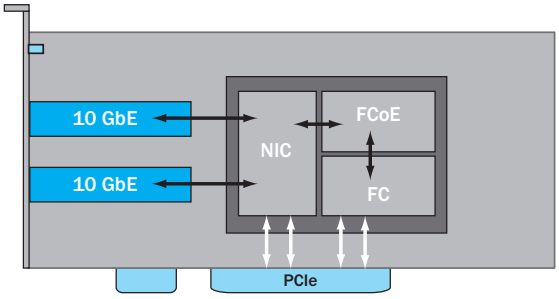

By avoiding changes to FC, FCoE CNAs are designed with HBA and NIC entities that perform the same functions found in today's HBAs and NICs. In addition, CNAs contain an FCoE entity that is internal to the adapter and not visible to the server Operating System (OS). Unlike HBA and NIC entities, which are exposed to the server, the FCoE entity is not exposed to the server nor is it seen by the server OS. In other words, when a CNA is installed in a server, the server OS will not see any FCoE devices. On the other hand, the server OS will see NIC and the HBA entities and expect device drivers for FC and Ethernet, so that the CNA can function as a storage and LAN interface adapter.

It is important to notice that as far as the server OS is concerned, nothing that is FCoE in nature is needed. In other words, FCoE has no impact on the server environment when it comes to device drivers or software architecture, and it does not require anything new from the user. It is also significant to note that FC and NIC drivers remain unchanged from today's environment. The result of preserving the existing server device drivers is that servers with CNAs can be added to today's FC environments without the need for any changes. IT professionals can continue to use the same FC drives they are familiar with and more importantly continue to use the same FC management applications they are accustomed to working with, such as Brocade® Data Center Fabric Manager (DCFM™).

Figure 4.
Server drivers from HBAs to NICs.

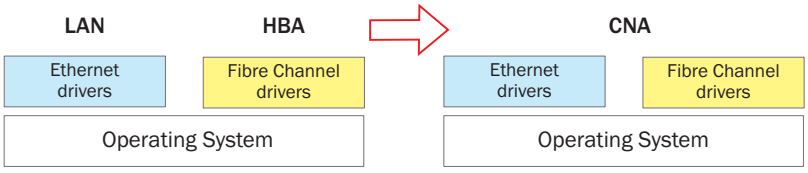

In short, CNAs, such as Brocade CNAs when they are available, are designed with a traditional FC HBA entity, a NIC entity where both are exposed to the server OS. CNAs also include an FCoE entity that is internal and does not alter the view of the server OS. The main role of the FCoE entity is to encapsulate FC traffic and send it to the CEE links and to do the reverse when traffic is flowing in the other direction. All traffic flows in and out of CNAs through eh 10 GbE CEE ports (see Figure 6).

There are three main functional block modules in an FCoE switch such as the Brocade 8000 Switch, as shown in Figure 5. The FC module in the Brocade 8000 is a fully functional FC switch supporting the features and services that are typically supported by Brocade FC switches. Connected to the FC switch is a bank of FC ports that are exposed to the outside for connecting to FC SANs. The Ethernet block of the FCoE switch is made up of a CEE switch supporting traditional 10 GbE Ethernet and CEE enhancement. A bank of 10 GbE CEE ports is exposed externally for connecting to corporate LANs. The CEE ports can be connected to

other CEE ports or to traditional 10 GbE ports. In the latter case, the CEE services on the CEE ports are downgraded as they communicate with traditional 10 GbE links.

FCoE switches contain an FCoE module or entity. As was the case in CNAs, the FCoE entity on the switch is also internal, is not visible to the outside, and has no external interfaces. The role of the FCoE entity is to extract FC payload form FCoE traffic and forward the FC frames to the FC switch on their way to FC storage targets. It also encapsulates FC traffic heading to the CEE switch to be transported back over CEE links.

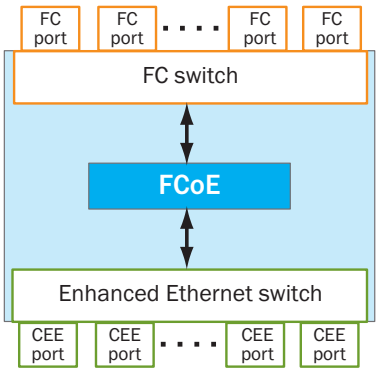

Figure 5.
FCoE switch architecture.

It is important to note that:

- Incoming traffic from the servers (CNAs) arrives at the switch over CEE links.
- Outgoing switch traffic leaves over FC ports to storage SANs and over CEE ports to LANs.

Upon receiving incoming traffic, the FCoE switch inspects the Ethertype of the incoming packets and uses it to determine the destination. If the Ethertype of the packet is FCoE, then the switch recognizes that the frame contains FC payload and forwards it to the FCoE entity. From there it is extracted and transmitted out over FFC ports. If the Ethertype is not FCoE, then the switch handles the traffic as CEE traffic and forwards it over CEE ports.

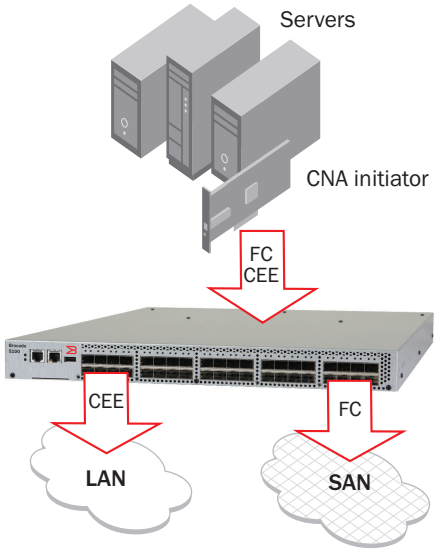

Figure 6.
FCoE traffic flows from server to switch and then out to the SAN and LAN.

In short, FCoE switches receive incoming traffic over CEE links, but send it out to intended destinations over FC ports to SANs and CEE ports to LANs. Since the FCoE entity in the FCoE switches is internal, FCoE switches can be seamlessly integrated into existing FC and Ethernet environments without the need for further modifications. Although nothing needs o be done to accommodate FCoE on the FC side, proper cables (CAT-6A) are needed on the CEE side.

BENEFITS OF SERVER I/O CONSOLIDATION

The benefits of FCoE and CEE will be most visible to data center professionals in consolidating server ports and the related cabling and switch ports. The promised savings will be realized when the deployment of the new technologies reaches a point where the economics of scale drive the cost of 10 GbE ports and related optics down to a price point that will accelerate adoption. Unless the cost of 10 GbE ports and optics becomes very attractive relative to technologies deployed today, the promised savings will remain as a promise and not reality.

Figure 7.
Server I/O consolidation..

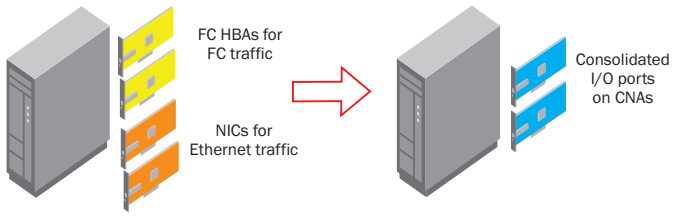

The positive impact of FCoE and CEE deployment will be reflected in many aspects of data center operations and management:

- **Cost of acquisition.** Reduced upfront capital outlays with FCoE and CEE, since they will be required to use fewer server CNAs, fewer cables, and fewer switch ports (switches) compared to technologies used today. The financial benefits will be magnified in the future when 10 GbE costs go down. However, it is important to note that initially FCoE and CEE will not be low-cost technologies.
- **Operating costs.** FCoE and CEE offer relief in in the area of power consumption and cooling, one of the highest expenditures in data centers, as fewer hardware components in each server will result in lower power consumption and cooling needs. The same applies on the switch side. Other operating costs that will start to decrease with FCoE and CEE are maintenance, SKU management costs, and other asset management expenses.
- **Decreased power.** Due to decrased server sprawl via a single adapter with the performance and capabilities of a separate HBA and NIC.
- **Other cost savings.** When FCoE is deployed, data centers will be able to continue to use existing FC management tools and not incur education or retraining costs. Less often cited are the time savings realized from dealing with simpler configurations and a much less cluttered environment—a result of reduced cabling and cable management needs. Troubleshooting and diagnostics can be performed more easily in environments in which technicians can identify and correct problems in shorter times. The reality is that simpler cabling helps reduce the potential for human errors.

FCoE USAGE MODELS

The benefits of FCoE and CEE deployment are best noted in three areas: consolidation using server virtualization and link consolidation, top-of-rack switch deployments, and end-of-row deployment using FCoE blades for backbone environments.

Consolidating Servers and I/O Links

Server virtualization is one of the drivers for higher I/O needs in servers. CNAs with 10 GbE CEE links provide virtualized applications with the IP links they expect and allow data centers to deploy multiple applications on individual servers. In addition to consolidating server I/O links over CEE ports, virtualization gives data centers the opportunity to consolidate applications over a smaller number of servers while benefiting from added application mobility. IT managers will also take advantage of the 10 GbE CEE links to consolidate slower Ethernet and FC connections onto the much faster connections offered by the new generations of CNAs. The combination of reducing server and I/O link counts results in noticeable savings and greater efficiencies in data centers.

Top-of-Rack Deployment

This deployment model offers the clearest example of how data centers can take advantage of the benefits of FCoE and CEE built into the Brocade 8000 Switch. The access layer is the most suitable area for this model of deployment either in a "greenfield" data center deployment or as additions to existing FC configurations. Today, most server racks are connected to company LANs using a pair of Ethernet switches for redundancy. The server racks are also connected to corporate SANs via redundant pairs of FC switches.

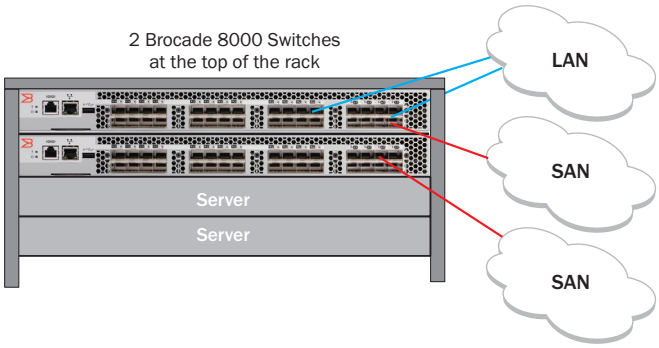

Figure 8.
Top-of-rack deployment.

Using the Brocade 8000, data centers can replace the Ethernet and FC switches with a single pair of FCoE switches that simultaneously connect the rack to corporate LANs and SANs. Savings in switch ports and rack space are evident, since rack servers are currently deployed with LAN and SAN separate connections and adapters. The Brocade 8000 CEE connections eliminate the needs for separate ports and cables, as both storage and LAN traffic travel over the same CEE cables utilizing a single server port and a single switch port for both types of traffic. The result is a reduced number of server adapters and a much simplified cabling inside the server cabinet.

End-of-Row Deployment

Driving FCoE and CEE capabilities to the backbone completes the reach of the new technologies from servers all the way to the core of the network. Brocade FCoE blades for the Brocade DCX® and DCX-4S Backbone make that reach a reality. Brocade FCoE blades deliver the features and capabilities found in traditional FCoE top of rack switches, but rely on other blades to provide the FC ports used for connecting SANs. With the FCoE blades, Brocade offers customers a complete solution that enables them to use Brocade DCFM framework to manage the organization resources from server initiators all the way to storage targets

SUMMARY

FCoE and CEE are two evolutionary technologies that help evolve the data center into a more efficient and simpler to manage environment. They also bring the industry closer to its goal of building greener data centers. The ability to transport SAN and LAN traffic over a single wire opens the door for true server I/O consolidation and reduced cabling clutter. FCoE deployments in the near and mid terms will preserve existing purpose-built SANs and LANs as FCoE switches continue to direct I/O traffic to its appropriate, but separate, destinations.

As FCoE is designed to encapsulate FC without altering its frames, FCoE solutions can be added to existing FC environments without any disruptions, thus preserving FC investments and extending their value. FCoE also allows customers to enhance the value of their FC management applications as they can extend their reach to FCoE solutions.

Brocade offers full FCoE and CEE solutions that extend from server CNAs to top-of -rack switches and end-of-row FCoE/CEE blades to a unified management tool that brings all of the elements together. With the new FCoE and CEE offerings, Brocade extends it data center industry leadership and continues its drive to transform the data center of the future

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