

Business white paper

FCoE and data center convergence

A business perspective



Table of contents

- 3** Introduction
- 3** Converging your data and storage networks
- 6** Technologies enabling convergence
- 7** Getting started with FCoE
- 8** Typical real-world deployments
- 9** Enabling the converged data center
- 10** Summing up

FCoE enables network I/O convergence

Converging your networked storage and data traffic onto a shared transport can simplify your networks and dramatically slash capital expenses and operating costs. Thanks to advancements in open industry standards and technology, Fibre Channel over Ethernet (FCoE) is emerging as the transport of choice for moving different data streams while preserving the unique advantages of each.

Introduction

“Cloud computing” is on everyone’s lips today. A central tenet of cloud computing is for dynamically allocated pools of resources to be available on demand for all deployed applications. Data centers are adding servers built to be more capable, with fast I/O and multi-core CPUs. Storage, too, is better and faster than ever. This infrastructure is being used more efficiently, thanks to virtualization techniques that create highly portable virtual machines and move them instantly between compute nodes and even distant data centers.

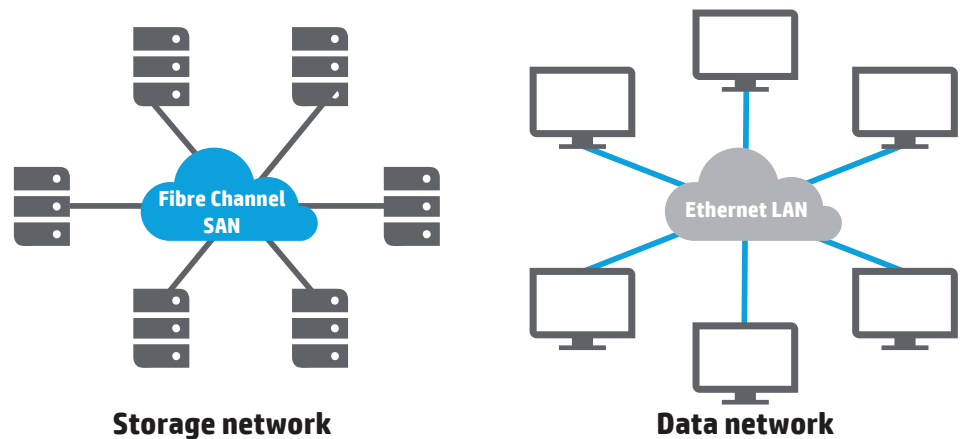
Yet for all the cloud-centric chatter about “infrastructure as a service,” one key factor is often overlooked: at the most fundamental level, cloud computing is still based on hardware. Often, it is the hardware I/O connects and infrastructure that limit an organization’s ability to take advantage of cloud computing and virtual machines. Adding increased server capabilities and virtual machines puts a premium on I/O ports, because those applications need links with adequate bandwidth for storage and data traffic.

Impelled by the need to keep pace with amped-up demands on infrastructure, more and more enterprises are moving to 10-Gigabit Ethernet (10 GbE) to carry the bulk of their network traffic. Faster network speeds mean better system response and greater productivity for end users, of course. However, the faster I/O speeds have an attendant, complementary benefit for the IT department: enabling them to converge data and storage traffic using FCoE technology.

Converging your data and storage networks

Behind the scenes of cloud computing, simplicity is the Holy Grail for your IT architects and administrators. By eliminating extraneous components and cables, you can save on hardware costs and also reduce management complexity.

A particular target for improvement is the current separation of storage and data networks. In fact, the deployment of a single data center fabric that carries networking, storage, and cluster I/O traffic is a goal that has been pursued by many companies using several technologies. Yet most data centers today still depend on different purpose-built networks—one for storage data and one for Internet Protocol (IP) data.

Figure 1. Most data centers rely on separate networks for storage data and IP data

Many enterprise data centers use Ethernet networks for LAN and IP data traffic, plus separate storage area networks (SANs) for Fibre Channel storage traffic. Here's why:

- The storage network needs to provide block I/O for applications such as network booting, mail servers, and large, data-intensive databases; and for this, the choice is usually Fibre Channel. Indeed, Fibre Channel SANs can help bring about storage consolidation and business continuance, while delivering centralized storage management, high performance, and reliability.
- Ethernet is the protocol for IP data. It's the method used by local area networks (LANs) to transfer 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 traffic, including corporate communications and files, voice-over-IP telephony, and even storage, using technologies such as NFS, CIFS (SMB), and iSCSI.

Technical differences

Storage traffic consists of block data, which does not tolerate dropped frames. That's why Fibre Channel is the basis for storage networks: it's lossless. In technical terms, Fibre Channel has a mechanism to keep track of buffer availability at the other end of the link. If the buffer is empty or only partially full, the sender forwards traffic, knowing that whatever frames it sends will be accepted by the receiver. If traffic is congested and the buffer is full, the receiver notifies the sender to wait. Therefore the receiver never needs to drop frames.


Ethernet, however, is a lossy network. It guarantees delivery by relying on the TCP protocol, which uses retransmission as the means of ensuring delivery. Ethernet switches don't signal to the sender that they are out of buffers. When switches are congested they tend to drop frames, which will be retransmitted later.

The downside of dual networks

Using two different means of data transfer actually works quite well, but it comes with a cost, especially in a blade server environment where space is at a premium. Ethernet IP data networks and Fibre Channel SAN fabrics 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.

Building and managing two separate networks for IP data and storage adds complexity and costs to the data center—hardly the simpler, more straightforward infrastructure that is the ideal. Thus, enterprises are looking for new ways to combine their IP and SAN networks. They want to enable the data center to run more efficiently and cost-effectively, while preserving their investments in both Ethernet LANs and Fibre Channel SANs. They're looking for convergence, with a single type of adapter in their servers, and both IP data and storage traffic sharing the same transport.

Figure 2. A converged network encompasses both data and storage, resulting in fewer components and easier administration.



Benefits at the server edge

The benefits of I/O convergence are most evident at the server edge. For uptime reasons, a server typically has two or more LAN connections and two or more Fibre Channel connections, and even more when higher performance is desired. On servers—blade servers in particular—the option slots occupied by Fibre Channel are a scarce resource, and the Fibre Channel host bus adapters (HBAs) themselves make up a noticeable portion of the server’s cost. Also, reducing the number of adapters helps to minimize heat dissipation in blade servers.

Network convergence enables you to reduce data center costs and complexity, and simplify network management, while protecting existing investments in Fibre Channel SANs. The benefits of convergence at the network edge are clear:

- **Lower capital expenses (CapEx)**—Because you purchase and operate a single shared transport instead of several, the savings can be substantial. Besides reducing the number of server adapters, you also save on cables, expensive transceivers, and switch ports. Using LOM (LAN on motherboard) Ethernet interconnects provides even more savings.
- **Reduced operating costs (OpEx)**—In a converged environment, the HBA and the NIC are effectively combined within a single device called a Converged Network Adapter (CNA). Reducing the number of server adapters and lowering the switch port count reduce the cost of power consumption and cooling, thus lowering operating costs.
- **Easier management**—With just one type of connection to each server, and fewer ports and switch connections, you reduce configuration and cabling errors. Management is substantially easier, too.
- **More efficient networking topologies**—The converged network is simpler, making deployment and operation more efficient.

Barriers to convergence adoption

The barriers to the adoption of convergence have been both technical and organizational. On the technical side, administrators were not eager to allow storage data to travel over the LAN, because they view Ethernet, with its dropped packets and frequent congestion, as being an unreliable transport mechanism. The technical challenges include:

- Converging traffic from two network types onto a single transport
- Ensuring data integrity for storage traffic
- Dealing with network congestion, where the TCP protocol relies on retransmissions
- Applying differentiated treatment or policies to different classes of traffic

What's more, data center managers often feel that using separate Fibre Channel and Ethernet networks works just fine, so why change?

The organizational challenges are more intractable. Large organizations usually have separate sets of dedicated administrators for managing IP networks and SAN traffic. Questions often arise regarding how to share the responsibility for managing the newly converged environment.

Converging your data center

It's becoming much easier to adopt convergence. In fact, a "perfect storm" of technology advancements has made the converged data center a reality. A major catalyst is FCoE, a technology that preserves the benefits of Fibre Channel while enabling the transport of Fibre Channel storage traffic over shared Ethernet links.


FCoE is gaining acceptance as an enabler of server I/O consolidation (that is, converged I/O in servers). In addition, the increased adoption of 10 GbE in the data center, combined with lossless 10 GbE technologies, makes it possible to transport Fibre Channel data flows with LAN IP data traffic on a shared Ethernet infrastructure.

Let's take a closer look at some of these enabling technologies.

Tech talk: How FCoE encapsulation works

FCoE encapsulates Fibre Channel frames within Ethernet frames.

Keys to encapsulation:



- The Fibre Channel payload is usually transported in a Fibre Channel frame, with a header and a cyclic redundancy check to ensure the data is correct.
- In FCoE, that Fibre Channel frame is encapsulated within an Ethernet frame, and is transmitted over the Ethernet network.
- At the destination node, the frame is unpacked, layer by layer, until the payload is accessible to the receiving application.

Note that FCoE requires jumbo Ethernet frames, which are larger than the standard 1.5 kB Ethernet frame size. Also, Ethernet must become lossless to carry storage data with integrity.

Technologies enabling convergence

Fibre Channel over Ethernet (FCoE)

In essence, FCoE is an encapsulation protocol that wraps Fibre Channel frames within Ethernet frames for transporting them over Ethernet links. Fibre Channel traffic and Ethernet traffic are transported on a shared Ethernet link, then directed to their respective networks. FCoE reduces the amount of hardware resources in the data center, while lowering power and cooling costs. There are also fewer support points when converging to a unified network, which helps reduce the management burden.

FCoE Initialization Protocol (FIP)

FIP is the FCoE control protocol responsible for establishing and maintaining Fibre Channel virtual links between pairs of devices. Its main purpose is to discover and initialize FCoE-capable entities connected to an Ethernet network.

Data Center Bridging (DCB)

Data Center Bridging is the key to making Ethernet lossless. DCB is an enhanced 10 GbE technology for use in data center environments. It provides several capabilities needed for convergence, including the elimination of data loss due to queue overflow (called lossless Ethernet), and the ability to allocate bandwidth on shared Ethernet links.

DCB is actually a collection of several independent IEEE standards:

- **802.1Qbb, Priority-based Flow Control (PFC)**—Allows identification and prioritization of traffic flows. During periods of heavy congestion, PFC ensures that congestion of one data flow doesn't affect the rest of the flows on the shared transport link. (Note that PFC is a single-hop mechanism and doesn't provide a solution for end-to-end network congestion.)
- **802.1Qaz, Enhanced Transmission Selection (ETS)**—Allows grouping of similar data flows and assigns identifiers to the created flow groups. ETS also gives different priorities to the created groups, and allocates a percentage of the available networking bandwidth of the shared transport link to each data flow group.
- **Data Center Bridging Exchange (DCBX)**—Extends the standard LLDP (Link Layer Discovery Protocol) to support resource discovery and connection establishment between nodes. DCBX allows both sides to exchange PFC and ETS attributes. It also allows exchange of application capabilities (including those of FCoE). DCBX works only for point-to-point interfaces.
- **802.1Qau, Quantized Congestion Notification (QCN)**—An end-to-end protocol that generates a congestion notification message (CNM) when a device experiences congestion. This tells the sending node to halt transmissions for a short period of time, and data traffic is throttled down until congestion is relieved. To be effective, QCN needs to be implemented by everyone on the network.

Getting started with FCoE

There's no need to rip and replace your old network to add FCoE to your environment. Your legacy network can easily coexist with new components as you build out your FCoE deployments. Remember, you aren't combining existing networks, you're just moving an increasing amount of storage traffic over shared Ethernet links.

For most organizations, we recommend a gradual FCoE deployment, one that preserves your investment in Fibre Channel SANs and storage as you wire in new converged equipment.

Start at the server edge

You'll find the greatest benefits of using FCoE when consolidating server I/O traffic—that is, for “server I/O consolidation” or “server edge convergence.” You'll probably want to start by implementing FCoE server I/O consolidation on blade servers or the top-of-rack (ToR) access layer.

What you'll need for convergence

- **Converged network adapter (CNA)**—An adapter that replaces a NIC and an HBA, combining both Fibre Channel HBA and Ethernet NIC functionality. To the operating system, the CNA appears as an adapter with the dual functionality of a Fibre Channel HBA and Ethernet NIC, so there's no need to change software stacks or card drivers.
- **Cables**—10 GbE connections from the CNA to the converged ToR switch. Within a rack, 10 GbE SFP+ direct attach cables are the most cost-effective.
- **FCoE switch**—Switches that support DCB and FCoE. You need these FCoE switches to connect your CNA-equipped servers to SANs and LANs.

Best applications for FCoE

If you're just beginning to implement FCoE, the best use is in applications where you can tolerate some risk—that is, in Tier-3 or Tier-2 applications such as SharePoint, email, research, and test and development. Save the transition to Tier-1 applications like billing, inventory, and especially online transaction processing until after your FCoE deployment is fully tested and proven.

Table 1. Best applications for FCoE

Application type	Fibre Channel over Ethernet (FCoE)		
	Separate Ethernet, Fibre Channel	Tier 2: Servers, business logic applications	Tier 3: Servers, Web access
Billing systems	•		
Inventory systems	•		
Research	•	•	
Email		•	•
Test and development		•	•
SharePoint			•

FCoE deployment considerations

The deployment of FCoE raises an obvious question: “Does this mean I can completely get rid of the separate Fibre Channel switches and the complexity of owning and operating them, and just have the one Ethernet network in my data center?”

Yes, you can do that today. However, as is the case with most technologies, there are tradeoffs to consider. FCoE makes the most sense for server I/O consolidation at the network edge. It is not as effective in the backbone. That's because most of FCoE's CapEx savings come from eliminating Fibre Channel interface cards (HBAs) in the servers, the Fibre Channel switch ports they are connected to, and the cables used to connect them. Once you've converged the

network edge using FCoE, there is little or no added CapEx benefit in replacing expensive Fibre Channel Director ports with an equal number of expensive Layer 3 Ethernet switch ports to get to an all-Ethernet data center.

Another question is, “What about end-to-end, multi-hop FCoE?” Successful implementation of FCoE through multiple hops in a data center network, or through an arbitrary network topology, requires technology that’s still evolving. Interoperability between multi-hop solutions from different vendors will also be a challenge for at least the next several years. Expect to apply significant technical resources configuring and tuning during early deployments of multi-hop FCoE environments.

Typical real-world deployments

Though there are some limitations, FCoE is a practical solution that is gaining ground in the enterprise. HP supports the deployment of FCoE at the server edge, and we can also provide end-to-end server convergence solutions, including HP Virtual Connect solutions for HP blade servers. We offer industry-leading servers, CNAs, LOM CNAs, and Ethernet switches with FCoE and DCB support. We also offer world-class storage and management, as well as server edge switches that support TRILL (Transparent Interconnection of Lots of Links), a new standard that replaces STP and enables multi-path L2 capability and multi-hop FCoE environments.

Some of our industry-leading server edge convergence solutions include:

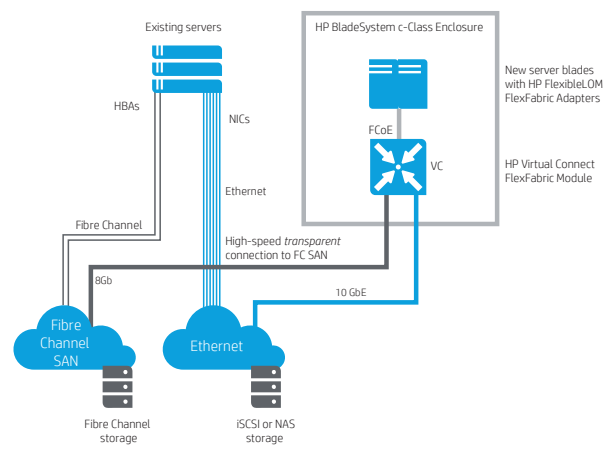
- HP 5820 10 GbE FCoE-capable ToR switch for top-of-rack deployment.
- HP 59x0AF Enterprise 10 GbE switch with DCB, FCoE, and TRILL support. The 59x0AF switch series also includes 10 GbE switches with deep packet buffers, and models with 10/100/1000BASE-T and 10GBASE-T ports.
- HP Virtual Connect blades for use with blade servers and connection to standard Ethernet or Fibre Channel switches. The HP Virtual Connect portfolio includes Ethernet, Fibre Channel, and converged fabric interconnect modules and firmware, giving you an ideal alternative to traditional switches and patch panels in HP BladeSystem c-Class enclosures.

Here are some real-world FCoE deployments that illustrate how you can take advantage of this innovative technology.

Deploying blade servers with FCoE and HP Virtual Connect

Figure 3 shows a simple, flexible way to connect virtualized server blades to data or storage networks. The HP Virtual Connect FlexFabric 10 Gb/24-port Module lets you add 10 GbE network and storage access to blade servers without disrupting your existing infrastructure. We expect 40 GbE links to be available in 2013.

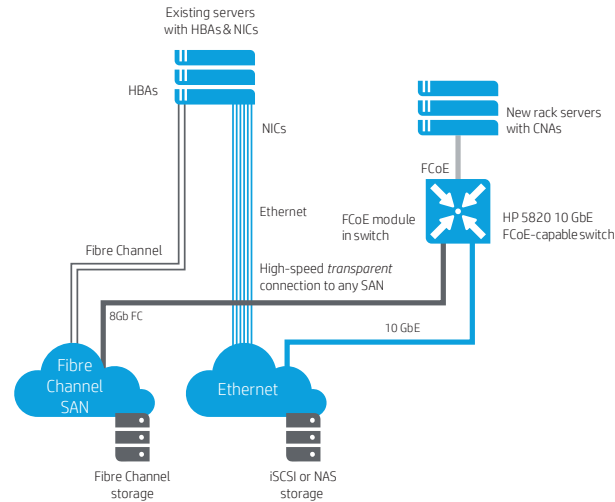
Figure 3. Adding blade servers and FCoE with HP Virtual Connect



Top-of-rack FCoE deployment

For a seamless and nondisruptive integration solution that provides investment protection, figure 4 illustrates how you can combine CNA-equipped rack servers with an FCoE-capable 10 GbE switch such as the HP 5820AF switch. This configuration allows the use of your existing infrastructure, while saving on capital outlays (fewer adapters) and operating expenses (less power, less heat, easier management) in your new equipment.


Figure 4. Deploying FCoE-capable HP switches and rack servers with CNAs



End-to-end FCoE deployment

Figure 5 illustrates an end-to-end, FCoE-ready deployment in existing environments. Based on industry standards, it includes an HP core switch with DCB support; HP 5820 10 GbE top-of-rack switches with Fibre Channel support; and converged HP 59x0 10 GbE ToR switches for server access with DCB, FCoE, and TRILL support. The HP 12500 core switch is FCoE/DCB-ready.

Figure 5. Deploying Fibre Channel, FCoE, and DCB



Enabling the converged data center

Fibre Channel over Ethernet offers plenty of advantages in terms of data center consolidation, lower capital and operating expenses, and easier management. However, since it's based on the Ethernet standard, it's not routable over IP networks, which means it is aimed at deployment in data centers. For wider deployments, innovative HP Intelligent Resilient Framework (IRF) technology and multi-hop technologies offer the potential to create massively scaled environments that eliminate complex, multi-tier topologies.

Contacts

Ahmad Zamer

Sr. Global Product Marketing Manager
ahmad.zamer@hp.com

Joe Zhao

Americas
joe.zhao2@hp.com

Olivier Vallois

EMEA
olivier.vallois@hp.com

Jason Xie

APJ
jason.xie@hp.com

Multiple hops with TRILL

The industry-standard Layer 2 networking encapsulation protocol TRILL provides shortest path frame routing and enables multi-hop environments. The protocol is designed to free available network bandwidth by replacing STP (Spanning Tree Protocol). TRILL also allows multi-hop links by utilizing routing bridges (RBridges) running the Intermediate System to Intermediate System (IS-IS) protocol. IS-IS is a link protocol that uses MAC addresses, eliminating the need for IP addresses.


TRILL provides multiple Layer 2 paths between Rbridges. It offers increased switching density, faster re-convergence times in response to configuration changes or failures, and improved network efficiency. TRILL lets administrators free up network bandwidth, and improves utilization by eliminating STP and opening all paths between nodes.

Flatter networks with HP IRF

HP IRF radically simplifies the architecture of networks, and delivers higher levels of performance while decreasing cost and complexity. HP IRF eliminates the need for an aggregation layer entirely. And, it works with TRILL, providing the perfect complement to this protocol.

The converged data center shown in figure 6 relies on core switches with IRF, DCB, and FCoE support. It also employs converged ToR access layer switches with IRF, DCB, TRILL, and FCoE links. We expect such Ethernet fabrics to be vendor-specific, at least in the near term.

Figure 6. The converged data center (2013 and beyond)



Summing up

Fibre Channel over Ethernet delivers the convergence of network traffic and FC storage today. It offers significant value to organizations looking to consolidate server I/O, network, and storage interconnects by converging onto a single network technology going forward. FCoE protects current investments in Fibre Channel networked storage, as it enables seamless integration with existing Fibre Channel SAN environments. Utilizing the same converged networking links for Fibre Channel storage and data networking traffic reduces your total cost of ownership, while preserving existing infrastructure investments and maintaining backward compatibility with familiar IT practices.

Learn more at
hp.com/go/networking

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document

