

Fibre Channel over Ethernet: A necessary infrastructure convergence

By Deni Connor, principal analyst

April 2008

Introduction

Consolidation of IT datacenter infrastructure is happening in all forms. IT administrators have consolidated servers with virtualization technology from VMware, Microsoft and Citrix and have reaped significant cost and operational benefits. They've consolidated their storage infrastructures with virtualization software and hardware. Now, IT is eyeing the consolidation of their network fabrics, concentrating on implementing a single converged fabric and dispensing with separate networks for LAN, SAN and high-performance inter-process communications (IPC) traffic. Convergence offers the promise of: fewer network adapters (HBAs, NICs and HCAs) installed in servers; reduced network cabling; fewer fabric switches (FC, Ethernet, and InfiniBand); reduced datacenter footprint; streamlined management; and, decreased energy consumption.

To make this I/O consolidation a reality, vendors have proposed enhancements to the ANSI IEEE 802 standard and a new specification to take advantage of these enhancements. These enhancements have been called Enhanced Ethernet (EE), Converged Enhanced Ethernet (CEE), DataCenter Ethernet (DCE) and Lossless Ethernet. Additionally, there is a new specification called Fibre Channel over Ethernet (FCoE), which is supported by companies like Cisco, IBM, EMC, QLogic, Sun, Emulex and Intel, that enables the Fibre Channel protocol to run natively over the IP-based Ethernet network in concert with other network traffic types. FCoE converges Fibre Channel SAN traffic with Enhanced Ethernet traffic, resulting in a single network that shares a common physical layer (fabric, cabling, transceivers and network interfaces). Ratification of the Fibre Channel over Ethernet standard is expected in 2009.

Simply put, Fibre Channel over Ethernet (FCoE) lets IT administrators run Fibre Channel traffic over a 10Gbit/sec. Ethernet network. Additionally, FCoE is appealing to IT, because IT departments are by nature conservative on spending and Fibre Channel over Ethernet lets them maintain their existing investments in Fibre Channel SAN technology.

The New Ethernet

Running Fibre Channel over Ethernet makes sense. Ethernet is ubiquitous in enterprise and SMB (small and medium business) networks. It is the protocol systems administrators are most familiar with. The advent of 10Gbit/sec. Ethernet and its enhancements are the catalyst toward I/O convergence. The combination offers the bandwidth and traffic management necessary to support multiple classes of traffic – storage, IP, IPC or management -- on a single cable.

There were some technical issues to overcome when merging Fibre Channel with 10Gbit/sec. Ethernet. Ethernet manages network congestion by dropping packets, which is fine for IP traffic, but not storage traffic, which requires a lossless traffic service. Storage traffic, which usually uses SCSI running over a Fibre Channel network, does not adapt well to lost frames and the latency a congested network causes. Enhancements to the IEEE 802 standard were necessary to allow Fibre Channel to run efficiently over Ethernet and make it lossless.

In order to dispel concerns of packet-loss behavior, congestion and latency, vendors who have a huge investment in Fibre Channel SANs have proposed enhancements to the Ethernet specification. These enhancements are what differentiate Enhanced Ethernet (EE) from standard Ethernet. Since EE solves the issue of lost packets, EE is also sometimes referred to as Lossless Ethernet.

Enhanced Ethernet (EE): Ethernet functions enable lossless networks

The proposed EE enhancements to the IEEE 802 standard are built on 10Gbit/sec. Ethernet and intended to provide lossless Layer 2 Ethernet. Four new functionalities are contained within EE: Priority Flow Control (PFC); Congestion Notification; DCB Capability Exchange Protocol; and, Enhanced Transmission Selection (ETS). Following is a definition of each of these enhancements:

Feature	Function	Benefit
Priority Flow Control(PFC)/Per priority PAUSE	Allows the network to selectively pause different classes of traffic	Allows for the creation of lossless lanes for FC while retaining packet drop congestion management for IP traffic
Congestion Notification	Throttles traffic at exterior of network in favor of interior traffic when congestion occurs	A lossless Ethernet in which packets are not dropped during congestion
DCB Capability Exchange Protocol	Allows Enhanced Ethernet devices to convey and configure their EE capabilities	Ensures a consistent configuration across the network
Enhanced Transmission Selection (ETS)	Allows lower priority traffic classes to use available bandwidth not being used by higher priority traffic classes	Maximizes the use of available bandwidth

These enhancements provide the ‘secret sauce’ that lets Fibre Channel run natively over Ethernet efficiently and reliably.

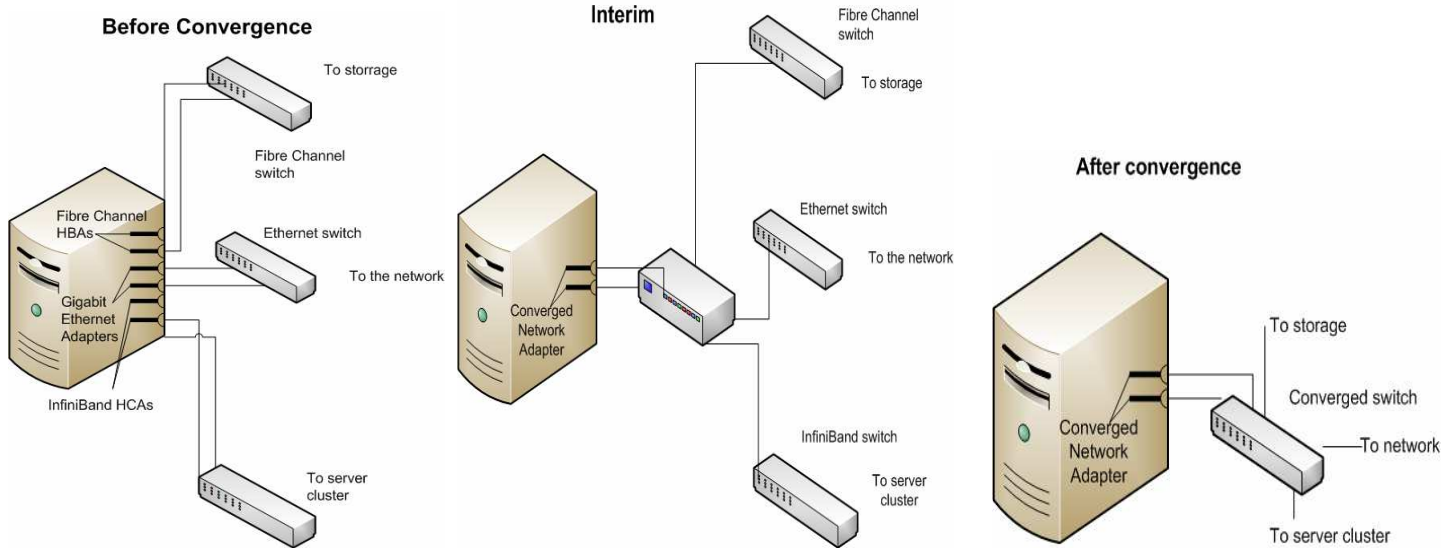
Encapsulating Fibre Channel into Ethernet

FCoE lets Fibre Channel traffic run over the enhanced 10Gbit/sec. Ethernet network without performance degradation or changes to Fibre Channel frames. Running Fibre Channel over Layer 2 Ethernet alleviates TCP overhead since Enhanced Ethernet provides a lossless environment.

In Fibre Channel over Ethernet (FCoE), complete Fibre Channel frames are encapsulated into Ethernet frames for transport over the Ethernet network. The encapsulation of the SCSI command set in the Fibre Channel frame does not require processing of TCP/IP overhead like iSCSI, Fibre Channel over IP (FCIP), and Internet Fibre Channel Protocol (iFCP).

Deploying Fibre Channel over Ethernet

Enterprise IT administrators looking to save on media, power consumption and floor space are expected to either adopt FCoE in a staged fashion, or migrate to complete convergence when they re-architect their datacenters. Visualize these implementation scenarios:



In the environment 'before' use of converged network adapters, as many as six adapters reside in each server – two redundant HBAs that connect to Fibre Channel switches, two redundant Gigabit Ethernet NICs that attach to the Ethernet switch, and two redundant HCAs for server clustering and attachment to an InfiniBand switch. Nothing has been replaced and no problems have been solved – there is still an over-abundance of media, too many network adapters and too many costly switches.

In migrating to a converged architecture, IT might simplify the network configuration and save on media costs and adapters by installing converged network adapters and a converged fabric bridge between the servers and Fibre Channel, Ethernet and InfiniBand switches. In doing so, they will not only reduce the number of adapters installed in the server, but also reduce the amount of media used to connect each segment by up to 66%.

Looking further out to when existing datacenters are re-architected for converged fabrics or when new datacenters are built, only one pair of redundant converged network adapters will reside in the servers, while the separate Fibre Channel, Gigabit Ethernet and InfiniBand switches will be replaced by converged network switches such as Cisco's Nexus 5000.

The QLogic QLE8000 Series – Leading the Industry in Converged Network Adapters

The QLogic QLE8000 series of converged network adapters have three distinct advantages over competing FCoE solutions. First, they offload all FCoE protocol processing from the server CPU and move it to the adapter. This provides superior performance versus competing FCoE NIC solutions that utilize software initiators running on the server CPU to process the FCoE protocol stack. QLogic's offload approach allows the CPU to do other tasks, such as running business applications. It also enables much higher levels of server virtualization by freeing the CPU to support more Virtual Machines (VMs).

A second QLogic advantage is that the QLE8000 series of CNAs will reuse the same battle-hardened, QLogic FC drivers that have already been deployed and proven in millions of existing datacenter servers. As QLogic is

the world market-share leader for Fibre Channel HBAs, their drivers have been developed and deployed for years and are considered the most reliable FC drivers in the industry. CNA vendors who are new to the Fibre Channel market, however, will need to develop Fibre Channel drivers from scratch and will need years of testing before they're 100% bulletproof and datacenter ready. Finally, QLogic is in the unique position of possessing proven Fibre Channel, Ethernet, and InfiniBand experience. Converged networking is meant to combine these three network types onto a single fabric that can carry data, storage, and inter-processor communication networking. Each of these networks has unique qualities and capabilities requiring different expertise to properly design and implement. As a market leader for Fibre Channel and iSCSI (which runs on Ethernet) HBAs and InfiniBand HCAs, QLogic is the only competing CNA vendor with experience in all three areas.

Advantages of Converged Fabrics

Fibre Channel over Ethernet and Enhanced Ethernet protect an organization's investment in existing networking hardware, software, and expertise by providing a bridge between existing Fibre Channel SANs and the Ethernet infrastructure. Installing converged network adapters reduces deployment costs and operational power and cooling requirements, as well as the number of personnel needed to support each separate fabric.

	Infrastructure Advantages
Preserve existing investments	Co-exist with Fibre Channel switches and HBAs; Ethernet switches and NICs
Common OS drivers and support	Uses same drivers as existing Fibre Channel HBAs (A QLogic advantage)
Cable efficiency	Reduces cabling by 66% by using a common media for storage, IP, and IPC
Green IT infrastructure	Consolidation of switches, HBAs, NICs and cabling reduces energy consumption
Rack space consolidation	Consolidating servers allows them to take less space in the datacenter, thus reducing the rack space footprint
Management consolidation	Consolidation of storage, IP and IPC networks allows management from a single converged interface
Maximized blade/multi-CPU servers	Converged I/O pipe allows performance and application improvements

SSG-NOW Assessment

Storage Strategies NOW sees FCoE being adopted in several market sectors: enterprise datacenters where the need for a converged network is required to support virtualized or blade server environments; and, in the small to mid-size enterprise, where the high end of the market will migrate their mix of Fibre Channel and IP SANs to Fibre Channel over Ethernet gradually to support business-critical applications. Other segments of the market such as small and mid-size businesses will continue to rely on iSCSI or network attached storage (NAS) deployments.

	Discrete Network	Converged Network		Discrete Network		
		1Gbit Ethernet	1Gbit Ethernet	10Gbit Ethernet (FCoE)	10Gbit Ethernet (FCoE)	8/16Gbit Fibre Channel
Fibre Channel/business-critical applications				✓	✓	✓
Enterprise datacenters			✓			
Small to mid-sized enterprises	✓	✓	✓			
Small to mid-sized businesses	✓	✓				

SSG-NOW estimates that uptake of the converged fabric will be slow in 2008, but ramp to over \$500 million as early as 2012.

Finally, it is important for customers to look to a vendor that can provide products for every market segment – from iSCSI, to 10Gbit/sec.Ethernet, to Fibre Channel to InfiniBand.

Note: The information and recommendations made by Storage Strategies NOW are based upon public information and sources and may also include personal opinions both of Storage Strategies NOW and others, all of which we believe to be accurate and reliable. As market conditions change however, and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Storage Strategies NOW, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.