

FIBRE CHANNEL OVER ETHERNET

A Review of FCoE Today

ABSTRACT

Fibre Channel over Ethernet (FcoE) is a storage networking option, based on industry standards. This white paper provides an overview of FCoE, describes the networking hardware and software components that makeup the ecosystem, and explains how the technology is deployed today.

October 2015

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, [contact](#) your local representative or authorized reseller, visit www.emc.com, or explore and compare products in the [EMC Store](#)

Copyright © 2015 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided "as is." EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

VMware and ESX are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number H5916.4

TABLE OF CONTENTS

EXECUTIVE SUMMARY4
 AUDIENCE..... 5

INTRODUCTION6

CREATION OF FCOE INFRASTRUCTURE.....6

FCOE HAS FIBRE CHANNEL (FC) INSIDE.....6

VIRTUALIZATION AND FCOE6

FCoE IS FC OVER 'LOSSLESS ETHERNET'7

ROAD TO A CONVERGED FABRIC.....8

CONCLUSION.....9

EXECUTIVE SUMMARY

The question for most data center managers is not should they network their storage, but which of many options available should they use and for which applications can they justify the expense. Using networked storage reduces not only the capital expenses by allowing centralized management and highly available configurations, but also operational expenses by allowing centralized management and simplified backup and replication. Server virtualization and IaaS increases the percentage of servers attached to networked storage in the data center.

Fibre Channel is the predominate choice for networked storage in the data center because it allows for the creation of a highly scalable – hundreds of thousands of nodes – and a reliable solution that can be centrally managed. For networking, servers typically have multiple Ethernet network interface cards (NICs) and Host Bus Adapters. The high number of server adapters, cables, and separate networking and storage fabrics adds to complexity and cost. As a result, many customers would like to be able to have all the applications run over a single converged network. This goal of I/O consolidation on a unified network for all traffic leads to a savings in infrastructure (cabling, reduced sparring, and so on) and a simplified management environment. Overall data center trends require solutions that take up less space, draw less power, and require less cooling; suppliers are reaching a breaking point where supporting multiple separate networks will not allow them to meet these requirements.

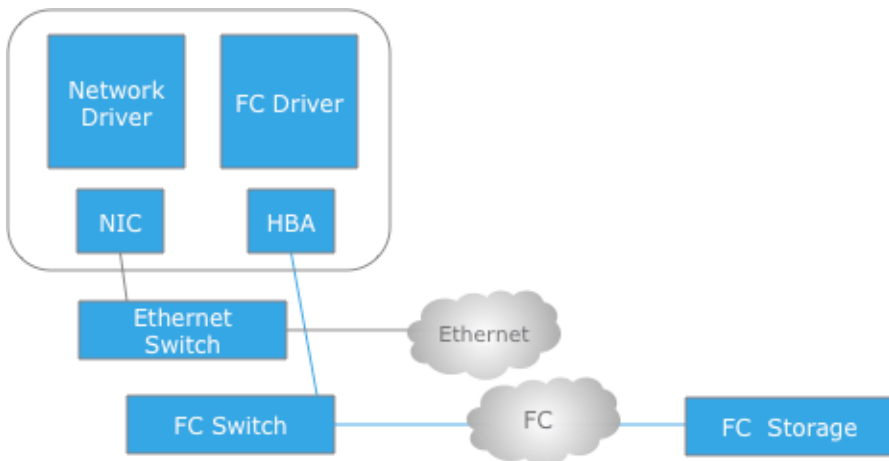


Figure 1 Existing data center with separate Ethernet and FC networks and multiple adapters

Ethernet has been given top considerations as a potential solution for I/O consolidation due to the large install base and broad general understanding of the technology by all IT departments. Fibre Channel over Ethernet (FCoE) allows an Ethernet-based SAN to be introduced into the FC-based data center without breaking existing administration tools or workflows. This is done by allowing Fibre Channel traffic to run over an Ethernet infrastructure. FCoE is simply encapsulating FC over the Ethernet portions of the connectivity, while the server applications as well as FC SAN and FC storage, the traffic is still FC. FCoE can be easily added into existing FC environments in the contrast to previous attempts at consolidation with iSCSI and InfiniBand, which required new drivers and rip-and-replace for customers that are using Fibre Channel.

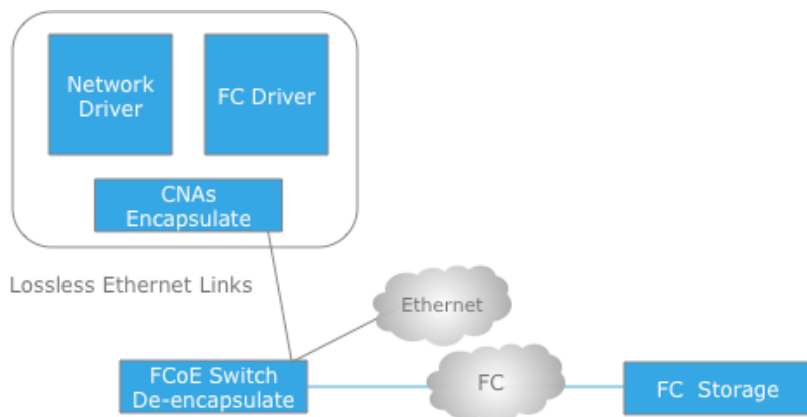


Figure 2 FCoE environment with converged adapters at the server attached to a converged switch.

AUDIENCE

This white paper is intended for network and storage administrators.

INTRODUCTION

Today's networks use different protocols to send information between devices. Traditional Ethernet is a family of frame-based computer networking technologies for local area networks (LANs), whereas Fibre Channel is used for storage area networking (SANs). Fibre Channel over Ethernet, or FCoE is a storage networking protocol that supports Fibre Channel natively over Ethernet. FCoE encapsulates Fibre Channel frames into Ethernet frames, allowing them to run alongside traditional Internet Protocol (IP) traffic. This white paper provides an overview of FCoE, describes the hardware and software components that make up the new ecosystem, and explains how the technology is deployed today.

CREATION OF FCoE INFRASTRUCTURE

FCoE requires the deployment of three new components: a Converged Network Adapter (CNA), Lossless Ethernet Links, and an FCoE switch. There are two types of CNAs available: a "hardware"- based solution – where the lower-level FC (exchange and sequence management) and FCoE functions are done in hardware – or a "software solution – where some of those functions are performed in software. When obtaining a hardware solution from a vendor that manufactures FC HBAs, the drivers will be the same for both FC and FCoE adapters; this provides the least disruption and easiest integration into existing environments. The software solutions require the development of the new FC drivers that will be run in software; one of the efforts can be found at Open-FCoE.org. The FCoE switch is the network device that will connect to existing LAN and SAN environments. The T11 standards FC-BB-5 group (<http://www.t11.org/fcoe>) created the FCoE protocol, which enables the creation of CNAs and FCoE switches. The link that FCoE traffic uses must provide the same level of lossless behavior that can be found with the Fibre Channel environments today. The Data Center Bridging group (part of the IEEE 802.1 standards, see <http://www.ieee802.org/1/pages/dcbridges/html>) has created an Ethernet environment that can support storage and the IPC traffic. This body of work is known as DCB Ethernet.

FCoE HAS FIBRE CHANNEL (FC) INSIDE

In creating the FCoE protocol, the goal is summed up in the name – take Fibre Channel (FC) and put it over an Ethernet infrastructure. To do this, the Fibre Channel frames are encapsulated – not translated or bridged – in an Ethernet frame (Figure 3). The mapping is 1:1, which means that there is no segmenting for FC frames nor are multiple FC frames put in a single Ethernet frame. It is a requirement that all devices for FCoE (adapters and switches along the path) support mini-jumbo frames to allow the largest FC frame to be supported without segmentation. Flow control of FCoE is controlled by Priority Flow Control mechanisms rather than by the buffer-to-buffer credits of FC. The naming conventions for FCoE port types are very similar to FC:

- N_Ports (HBAs and storage) are VN_Ports (CNAs or FCoE storage).
- F_Ports (fabric ports) are VF_Ports (FCoE switch ports that attach to VN_Ports).
- E_Ports (switch to switch) are VE_Ports (ports between two FCoE Ethernet switches).

Virtualization and FCoE

Currently, server virtualization environments (including VMware® and Microsoft Hyper-V) can choose among FC, iSCSI, and NAS for networked storage. The hypervisor has a storage driver stack that presents FC (from an HBA) and iSCSI (from a NIC) traffic to individual virtual machines (VMs) or guests as storage. There is also a virtual switch in the hypervisor that sends traditional LAN traffic to the VM/guest as network traffic. FCoE solutions using a CNA will function equivalently as existing solutions – the CNA will convert FCoE traffic to FC packets in the hardware, so the hypervisor and VM/guest will still work as if physical FC HBAs and NICs were installed. Software FCoE solutions with a standard NIC will require additional developments before they can be used. Currently hypervisors will not identify FCoE traffic that is still encapsulated and virtual switches (including Cisco Nexus 1000V in VMware ESX®) do not have Lossless Ethernet functionality, so sending traffic directly to the VM/guest would not be reliable. FCoE will expand the available storage networking solutions options for the high-growth server virtualization market. Moving to a converged fabric will simplify mobility by reducing the number of networks that need to be moved. It also allows a standard of configuration to be built, allowing for flexible storage deployment.

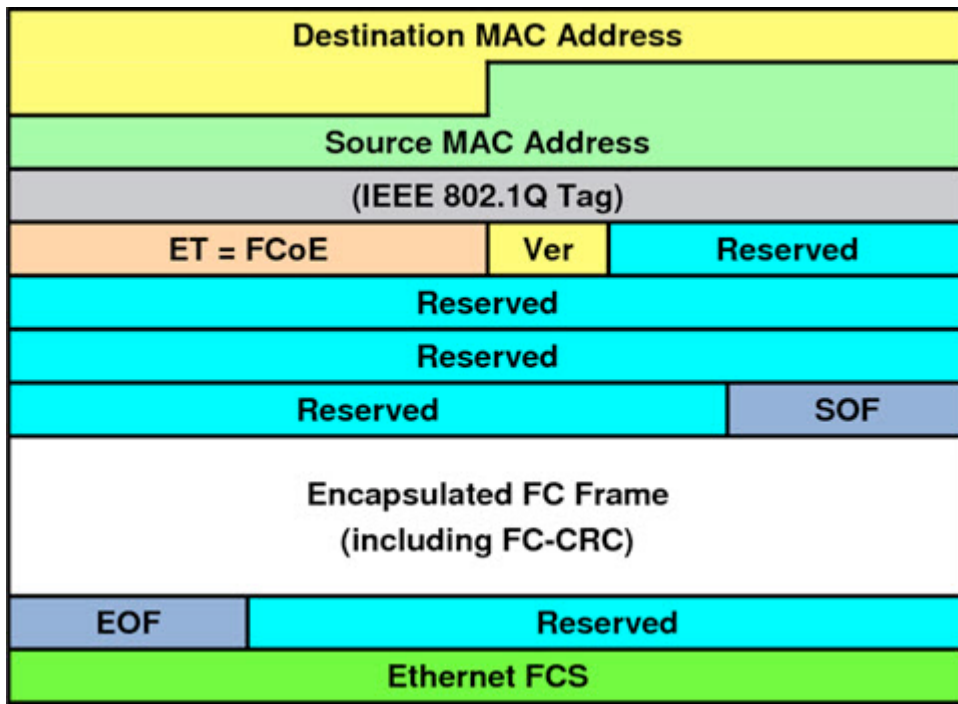


Figure 3 FCoE Frame Format

In Fibre Channel, security is typically not high on the list of discussion points because a FC SAN is isolated from the LAN so, many security problems are implicitly addressed. In the early days of FC, the small-scale physical security and optical cabling were usually sufficient for customers. The Fibre Channel Security Protocols standard (FC-SP) addresses how to protect against security breaches. Ethernet has its own well-established security practices, so the standards were only concerned with specific threats that are not covered in a typical Ethernet environment. An issue of concern was that while FC links are always private (two ends), FCoE links can be shared (three-plus ends) due to the nature of Layer 2 Ethernet. Access Control Lists (ACLs) will be used to specify on a port basis what operations are allowed to be performed, similar to the protection provided by a firewall. A feature called "Dynamic ACLs" is supported by the FCoE switches to reduce the administrative burden of manually configuring ACLs. FC-BB-5, the T11 standard for FCoE was ratified in June 2009.

FCoE IS FC OVER 'LOSSLESS ETHERNET'

The Ethernet infrastructure, over which FCoE will travel, must be of the lossless nature. Since the FCoE protocol does not contain TCP, any lost packets would require recovery at the SCSI layer. In a typical environment, based on existing Ethernet, this would happen much too often to be useable for storage environments. Fortunately, the Ethernet community has been looking at the issue of creating a lossless environment for a number of years. The set of features required to support lossless behavior is referred to as Data Center Bridging (DCB) and is being worked on by the IEEE DCB working group. The first set of features is link level protocols and by definition only covers the link between two devices (that is, either an adapter and switch port or two switch ports); they are Priority Flow Control (PFC) and Enhanced Transmission Selection (ETS).

Even before FCoE was conceived, there was a PAUSE function (Link Pause, IEEE 802.3 Annex 31B) that would allow for the creation of a lossless environment by stopping all traffic on a port when a full queue condition was achieved. The problem with classical PAUSE is that you cannot specify which traffic should be lossless, and as a result there is no mechanism to deal with the congestion created on the classes of traffic when a PAUSE is issued. Priority Flow Control (IEEE 802.1Qbb) creates new PAUSE function that can halt traffic according to the priority tag while allowing traffic at other priority levels to continue. Administrators use the eight lanes defined in IEEE 802.1p to create virtual lossless lanes for traffic classes like storage (that require lossless behavior) and lossy lanes for other classes.

In order to make sure that one traffic stream does not take too much of the overall bandwidth, the addition of Enhanced Transmission Selection (IEEE 802.1Qaz) is important for creating an environment storage and other traffic can share the same link. A common management framework can be created for bandwidth management. High-priority traffic like storage can be prioritized and allocated bandwidth while still allowing other traffic classes to utilize the bandwidth when it is not being fully utilized. The Data Center Bridging Exchange Protocol (DCBX) is responsible for the configuration of link parameters for Data Center Bridging functions. It determines which devices support the enhanced functionalities that create the "DCB" cloud where FCoE traffic can safely travel. It also allows for transparent pass-through to non-DCB traffic classes.

These link level enhancements start to create a safe environment for FCoE traffic. PFC, ETS, and DCBX have all been defined therefore solutions that have these functions shipping today should be able to support FCoE.

ROAD TO A CONVERGED FABRIC

Many customers are already running their storage traffic over Ethernet by utilizing iSCSI or NAS. FCoE is not designed to replace these solutions. In general, iSCSI environments tend to be small configurations with under 100 servers, while most FC customers are scaling into hundreds or thousands of nodes. FCoE can be plugged into these existing environments and 10 Gigabit Ethernet will provide the common infrastructure, that can support FCoE, iSCSI and NAS (Figure 4) and customers will choose based upon their requirements and skill sets.

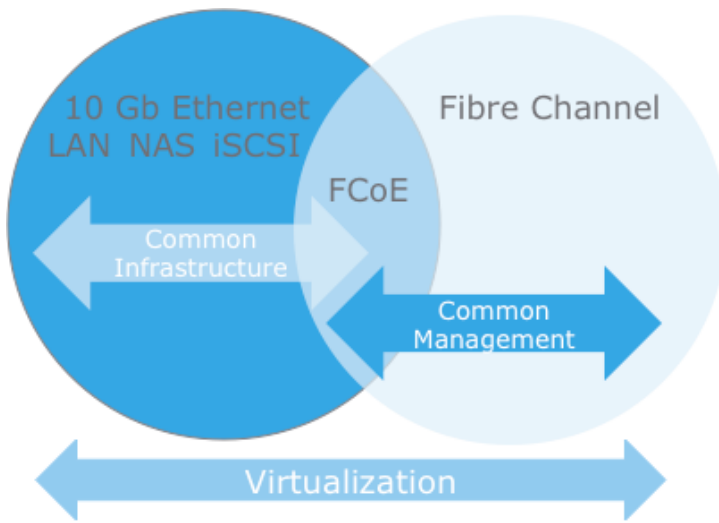


Figure 4: 10GbE Common Infrastructure for Storage

Today FCoE supports convergence at the server, in which the separate network (NIC) and storage (HBA) interfaces on the server are replaced by the CNA that is directly attached to the FCoE switch. A new rack of servers can deploy these technologies, while the existing storage and network environments remain unchanged. FCoE can also support convergence at the storage port and on switch-to-switch links, allowing for an end-to-end (FCoE multi-hop) configuration to be created.

No discussion of FCoE is complete without addressing the cabling infrastructure supporting the solution. When customers deploy a physical cabling infrastructure, they are looking for solutions that can be used for five to ten years. For 1 Gigabit Ethernet, the primary options for cabling are copper (1000Base-T with RJ-45 connectors) and optical (same physical cabling as FC); copper dominates this market with billions of ports installed while optical has historically been 1 percent of Ethernet ports. A copper option known as Twinax has become available for FCoE solutions of 10 Gigabit Ethernet. The option is based on the SFF-8431 standard and uses the SFP+ interface for a copper connection that is low cost and low power. The Twinax solution is limited to short distances, supporting between 1-10 meters, which is sufficient for server to top-of-rack or end-of-row switch environments. Standard multimode optical cabling will be used for environments that require longer distances such as from the rack to the core.

CONCLUSION

FCoE is a natural extension of FC networks that allows data centers to leverage converged I/O and extend the operational efficiencies of networked storage to a new wave of servers. Today, FCoE is most often used between the server and the ToR switch. At this time, end-to-end FCoE has not been widely deployed. That said, sometimes it takes years for technology like FCoE to be widely adopted.