

Small Computer System Interface over IP and Fibre Channel over Ethernet: A Comparison

Introduction

Data centers typically run multiple separate networks, including an Ethernet network for client-to-server and server-to-server communications, and a Fibre Channel SAN. To support various types of networks, data centers use separate, redundant interface modules for each network, such as Ethernet network interface cards (NICs) and Fibre Channel interfaces in their servers, and redundant pairs of switches at each layer in the network architecture. Use of parallel infrastructures increases capital costs, makes data center management more difficult, and decreases business flexibility.

A unified fabric can meet these challenges, consolidating I/O in the data center and allowing Fibre Channel and Ethernet networks to share a single, integrated infrastructure (Figure 1). Fibre Channel over Ethernet (FCoE) is one of the major components of a Unified Fabric. FCoE is a new technology developed by Cisco that is standardized in the Fibre Channel Backbone 5 (FC-BB-5) working group of Technical Committee T11 of the International Committee for Information Technology Standards (INCITS). Most large data centers have huge installed bases of Fibre Channel and want a technology that maintains the Fibre Channel model. FCoE assumes a lossless Ethernet, in which frames are never dropped (as in Fibre Channel) and that therefore does not use IP and TCP.

Small Computer System Interface over IP (iSCSI) is a mature technology created by the Internet Engineering Task Force (IETF) in the IP Storage (IPS) working group. It is based on the IP protocol stacks, it assumes (as IP does) an underlying unreliable network, and it delegates to TCP the recovery of lost packets.

iSCSI and FCoE are two important technologies that Cisco supports to provide Unified Fabric solutions.

Figure 1. Customer Benefits of Unified Fabric: Fewer NICs, Host Bus Adapters (HBAs), and Cables and Lower Capital Expenditures and Operating Expenses



iSCSI and FCOE Comparison

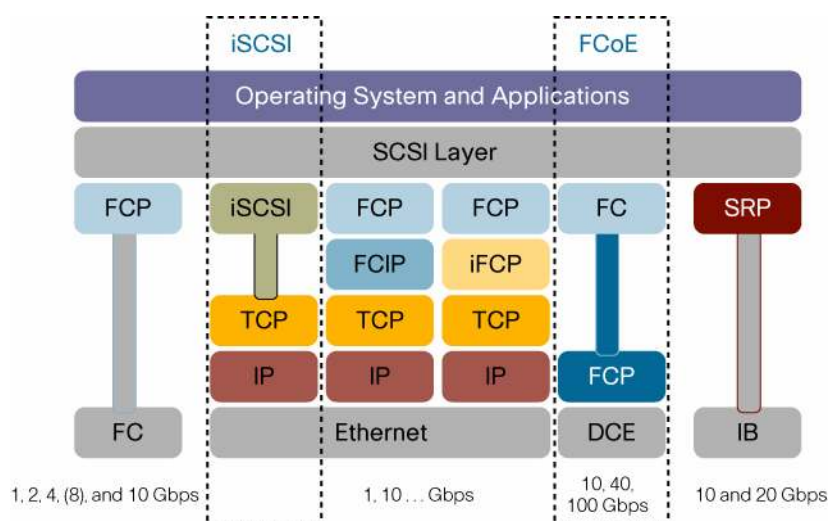
The value proposition of Unified Fabric is well understood: It dramatically reduces the number of adapters on the servers, the number of switch ports, and the associated cabling; it improves airflow; and it decreases capital and operating expenditures. Unified Fabric leverages IEEE Data Center Bridging standards to provide a lossless 10 Gigabit Ethernet network infrastructure that benefit both Layer 2 and 3 transports and are therefore valuable independent of the I/O consolidation technique used: iSCSI, FCoE, network attached storage (NAS), or, most likely, some combination of these.

The major server vendors offer iSCSI as a connectivity option, to provide virtual machines with block-level access to storage volumes, without the need to deploy high-performance Fibre Channel hardware. iSCSI has been proven capable of supporting enterprise-class applications assuming that the specific solution can reach the desired level of

performance and scalability and that management is simple enough not to negatively affect the total cost of ownership (TCO).

Figure 2 shows a comparison of the protocol stacks.

Figure 2. iSCSI and FCoE Protocol Stack Comparison



iSCSI fits well in the small and medium-sized business (SMB) market, where often price is more a concern than performance. iSCSI can use older Ethernet switches, whereas FCoE requires lossless switches. Until now, iSCSI has been limited to low-performance servers, mainly because Ethernet had a maximum speed of 1 Gbps, while Fibre Channel host bus adapters (HBAs) had 2- and 4-Gbps interfaces.

This performance factor is even more relevant for iSCSI storage arrays, in which a Gigabit Ethernet interface is typically shared by multiple servers to be cost effective. To overcome the bottleneck from back-end array connections, iSCSI gateways of up to 4-Gbps Fibre Channel have been deployed.

Although 10 Gigabit Ethernet has removed this limitation, there are concerns that the TCP termination required by iSCSI is onerous at this speed.

For some enterprise customers with a large installed base of Fibre Channel, the downside of iSCSI is that it is “SCSI over TCP,” not “FC over TCP,” and therefore it does not preserve the Fibre Channel management and deployment model. It has a different naming scheme (perhaps a better one, but anyhow different), different zoning, etc.

In contrast, FCoE integrates transparently into an existing Fibre Channel environment. FCoE is simple, and it contains the minimum indispensable information to carry Fibre Channel over Ethernet, and nothing else. IP and TCP are not included, and therefore FCoE is not IP routable.

Table 1 compares the main benefits of FCoE and iSCSI.

Table 1. FCoE and iSCSI Benefits Comparison

FCoE	Benefits
<ul style="list-style-type: none"> Mapping of Fibre Channel frames over Ethernet Fibre Channel enabled to run on a lossless Ethernet network 	<ul style="list-style-type: none"> Wire server only once Fewer cables and adapters Software provisioning of I/O Interoperates with existing Fibre Channel SANs No gateway; stateless




iSCSI	Benefits
<ul style="list-style-type: none"> • SCSI transport protocol that operates over TCP • Encapsulation of SCSI command descriptor blocks and data in TCP/IP byte streams  <p>The diagram shows a protocol stack. At the top is a purple box labeled 'Operating System and Applications'. Below it is a grey box labeled 'SCSI Layer'. Below that is a green box labeled 'iSCSI'. Below that is a yellow box labeled 'TCP'. Below that is a red box labeled 'IP'. At the bottom is a grey box labeled 'Ethernet'.</p>	<ul style="list-style-type: none"> • Wire server only once • Fewer cables and adaptors • New operational model • Broad industry support; OS vendors support their iSCSI drivers, gateways (routers, bridges), and native iSCSI storage arrays

Table 2 compares the positioning of FCoE and iSCSI.

Table 2. FCoE and iSCSI Positioning Comparison

FCoE	iSCSI
<ul style="list-style-type: none"> • Targeted for enterprise and service provider data centers • Suitable for data center applications • Use same management model as Fibre Channel • Has minimal effect on operations 	<ul style="list-style-type: none"> • Applicable for smaller environments (small and medium-sized businesses [SMBs]) • Suitable for desktop • Uses new management model • Requires stateful gateway
FCoE and iSCSI both run better over Unified Fabric	

The industry has already developed other standards to carry Fibre Channel over IP: FCIP. FCIP is routable, and it is already part of the FC-BB-5 standard that will also include FCoE.

Figure 3 shows a disaster recovery solution in which two data centers that use FCoE are connected to an IP network using FCIP. The IP network can be implemented with any available technology and does not need to be lossless.

Figure 3. FCoE and FCIP



In addition FCoE supports SAN booting as well as Fibre Channel does and definitively better than iSCSI.

A big change that is happening in the data center is the adoption of virtualization (for example, VMware and Xen). Virtualization is a means of consolidating multiple underutilized logical servers on few physical servers. Virtualization dramatically increases the I/O requirements, which may create another opportunity for Fibre Channel and FCoE since they do not require the TCP stack and therefore are not subject to TCP/IP overhead.

With virtualization comes virtual machine mobility, which often requires Layer 2 connectivity. This requirement fits well with Unified Fabric as it provides better congestion and bandwidth management ,and is therefore also synergistic with FCoE.

In practice, whether FCoE is faster than iSCSI is going to be as much a function of the particular implementation of the technology as anything inherent in the protocols themselves.

In terms of storage arrays, native iSCSI storage arrays have existed from some time, and native FCoE storage arrays have been shown by different vendors and will become a reality in 2009.

Conclusion

- FCoE was not designed to make iSCSI obsolete. iSCSI has many applications that FCoE does not cover, in particular in low-end systems and in small, remote branch offices, where IP connectivity is of paramount importance.
- Some customers have limited I/O requirements in the 100-Mbps range, and iSCSI is just the right solution for them. This is why iSCSI has taken off and is so successful in the SMB market: it is cheap, and it gets the job done.
- Large enterprises are adopting virtualization, have much higher I/O requirements, and want to preserve their investments and training in Fibre Channel. For them, FCoE is probably a better solution.
- Customers who have adopted Cisco® MDS 9000 family switches will probably prefer FCoE as it offers inherent coexistence with Fibre Channel, with no need to migrate existing Fibre Channel infrastructures.
- FCoE will take a large share of the SAN market. It will not make iSCSI obsolete, but it will reduce its potential market.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCSI, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0903R)