

Converging SAN and LAN Infrastructure with Fibre Channel over Ethernet for Efficient, Cost-Effective Data Centers



OVERVIEW

Fibre Channel over Ethernet (FCoE) is an emerging data transport technology that simplifies the network fabric by converging both storage area network (SAN) and LAN connectivity onto a single 10 Gigabit Ethernet link. As the use of SANs in the data center continues to grow, FCoE can decrease the cost of connecting servers to the SAN, while also preserving the existing management infrastructure and reducing overall power consumption.

By encapsulating Fibre Channel payloads in Ethernet frames, FCoE uses a switch that supports the technology, such as Cisco® Nexus 5000 Series Switches, to connect transparently to existing environments. As a result, organizations can simplify their infrastructures, with fewer server network interfaces, reduced cabling infrastructure, and a unified switching architecture, which also reduces power requirements in the server room. At the same time, FCoE facilitates connection of a higher percentage

of servers to the SAN, taking better advantage of SAN resources overall. SAN traffic running on FCoE can provide the same latency, security, and traffic management characteristics as if it were running on Fibre Channel.

A set of network technologies known collectively as Data Center Bridging (DCB) enables Ethernet fabrics to support lossless transmission, making them suitable for carrying SAN traffic. This set of emerging standards enables better traffic prioritization over a single interface, as well as advanced means for shaping traffic on the network to decrease congestion.

Compared to separate LAN and SAN interfaces provisioned to connect servers to the SAN fabric, a network using FCoE can be simpler and less costly to operate, requiring less power and equipment, making it more environmentally sound.

Evolving SAN Environments

Deployment of SANs using both the common connection fabrics—Small Computer System Interface over IP (iSCSI) and Fibre Channel—is increasing dramatically, and that growth is expected to continue through the foreseeable future, as shown in Figure 1. Growth in deployment of iSCSI-based SANs is especially prevalent in medium-sized businesses, branch offices, and new installations, while Fibre Channel SANs remain the choice for enterprise-scale deployments.

SAN Growth: Fibre Channel and iSCSI Storage Systems (in Petabytes)




Figure 1. Deployments of Fibre Channel and iSCSI SAN systems are expected to continue growing at least through 2011. (Source: Worldwide Disk Storage Systems 2007–2011, Forecast Update, IDC, No. 209490, December 2007)

Although use of Ethernet-based storage with iSCSI is growing quickly, most installed SANs today are Fibre Channel, and as IT managers deploy new servers in data centers, they need to be able to access the existing storage networks. FCoE lets them do this inexpensively and easily.

A number of factors promote this growth. Remote replication of SAN data enables disaster recovery and allows data to be centrally managed but accessed from multiple sites on a local basis. Point-in-time snapshots are also a widely used SAN capability, allowing organizations to maintain backups and audit trails that are helpful operationally and that also facilitate regulatory compliance. SANs are also typically needed for live migration of virtual machines for automatic failover and load balancing between servers. SANs also enable storage resources themselves to be virtualized, decoupling the data-storage entities from the physical hardware used to store them.

In addition to advanced features, SANs provide increased availability, which is increasingly important as more applications throughout the enterprise come to depend upon access to stored data. SANs also decrease the need for empty hard drive space on individual servers to accommodate future growth. Instead, extra storage space can readily be added as needed to a centralized point by means of a SAN topology.

For the servers in the environment to take advantage of benefits like these, they must be connected to the SAN. With Fibre Channel SANs, one or more host bus adapters (HBAs) must be purchased for each server, which adds considerably to equipment costs. For mission-critical applications (and often others), most organizations provide redundant connectivity, driving costs even higher. It is necessary with Fibre Channel SANs to operate separate networks for the LAN and SAN environments, as shown in Figure 2. These separate networks require added expense due to requirements such as increased numbers of network interfaces, additional cabling and switch ports, and more complex support requirements. Those expenses become even greater as the environment grows over time.




Figure 2. Today, most data centers use separate LAN and SAN networks with separate switches and network adapters in the server; for Fibre Channel SANs. This topology requires a dedicated Fibre Channel infrastructure.



Figure 3. In FCoE, the Fibre Channel payload is encapsulated in Ethernet frames.

Unified Networking with Fibre Channel over Ethernet

It is now possible to carry both LAN and SAN traffic over a single Ethernet network. FCoE extends Fibre Channel traffic onto a lossless 10 Gigabit Ethernet fabric, converging LAN and SAN I/O onto one set of cables. By means of an FCoE-capable switch, this technology connects transparently to existing Fibre Channel networks, coexisting with the rest of the topology. That coexistence enables organizations to implement FCoE incrementally, reducing effort and risk.

Transmission of Fibre Channel traffic over Ethernet fabric requires the encapsulation of native Fibre Channel frames into Ethernet packets, as shown in Figure 3. This methodology preserves the native format, so FCoE traffic appears as Fibre Channel traffic to the Fibre Channel fabric. This characteristic allows IT organizations to maintain the existing environment’s latency, security, and traffic-management attributes before and after migration. It also preserves investments in Fibre Channel expertise and equipment.

For more information about how this encapsulation is structured and how the data is transmitted using FCoE, see “Fibre Channel over Ethernet in the Data Center: An Introduction” at http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns783/white_paper_FCIACoE.pdf. FCoE is a standards-

based technology for storage networking; ANSI T11 FC-BB-5 was adopted in June 2007, and ratification is expected in 2008, with strong support from the storage industry as a whole.

The main features of FCoE and a unified fabric are summarized here.

Simplified Infrastructure

A unified data center fabric that incorporates FCoE reduces overall infrastructure requirements by enabling new servers to access the SAN without connecting to the Fibre Channel network. The unified fabric therefore enables IT to install and maintain fewer network adapters, switch ports, and cables.

Table 1 summarizes these advantages for a hypothetical network of 16 servers that uses redundant network interfaces for SAN and LAN connectivity in comparison to a network that implements FCoE end to end.

In addition to decreased requirements for the cables themselves, the simpler topology helps reduce cabling errors and make day-to-day tasks in the server room easier, allowing new servers and racks to be provisioned more quickly. Reducing cable clutter also helps avoid restriction of front-to-back airflow, which can help improve cooling efficiency. Intel estimates that implementing FCoE can save up to US\$2600 per server.¹

Table 1. Requirements for separate SAN and LAN networks compared to unified fabric.

	Requirements for Separate SAN & LAN			Requirements for Unified Fabric		
	10GbE (no FCoE)	Fibre Channel	Total	10GbE (FCoE)	Fibre Channel	Total
Network Adapters/ Ports	16/32	16/32	32/64	16/32	0	16/32
Switches	2	2	4	2	0	2
Cables	32	32	64	32	0	32

Logical Fit with Rest of Hardware Upgrade Lifecycle

Implementing FCoE in an IT environment does not require changes to the core network. Support for FCoE traffic will typically require an upgrade to one or more edge switches, such as embedded blade and top-of-rack switches, and does not affect the core switching equipment or topology. Moreover, this switch upgrade is only an incremental addition to the upgrade to 10 Gigabit Ethernet from Gigabit Ethernet that many organizations will be undertaking in the next year or two. More powerful multicore servers support higher workload levels than previous machines, which in turn require the greater network throughput of 10 Gigabit Ethernet.

The adoption of FCoE is a logical addition to mainstream network design strategies. Because FCoE is compatible with existing Fibre Channel topologies, new systems can be deployed using FCoE side by side with the existing network environment. This strategy means less disruption to the organization as a whole, as well as easier integration into operations and budgets.

Preservation of Existing Management Infrastructure

Organizations that already have Fibre Channel based SANs in place can use their existing management infrastructure, protecting their investment in management applications, expertise, and training as well as simplifying implementation. Because FCoE uses the same protocols as traditional Fibre Channel, the management framework is the same, regardless of whether the environment is based on traditional Fibre Channel, FCoE, or a combination of the two. Only the lowest two layers of the five-layer Fibre Channel model change.

Reduced Power Consumption

The smaller number of network interfaces and switches used with FCoE can reduce the power requirements in server rooms substantially. For example, each Fibre Channel HBA may consume about 12.5 watts, and a typical Fibre Channel switch may consume about 70 watts. Cooling the server environment requires additional energy equal to approximately 0.8X to 1.5X the input power.² For a typical rack, the power savings may be 400 watts from removal of the two HBAs from each of 16 servers,

plus an additional 140 watts from elimination of the two switches. Multiplying this number by approximately 2 to account for cooling results in a savings of roughly 1080 watts per rack. For a medium-sized or large enterprise, this combination of factors can represent a significant power savings. In addition to the potential for cost savings, every watt of power that an IT infrastructure conserves has a positive net effect on the environment. As companies look for new ways of making their operations more “green,” such discoveries are welcome.

No-Drop Data Center Bridging

Because conventional Ethernet is a best-effort topology, it drops packets in response to traffic congestion, which makes it unsuitable for use in storage environments. SANs typically use Fibre Channel or iSCSI to overcome this limitation, adding cost and complexity to the environment. FCoE provides a number of mechanisms that contribute to no-drop behavior over an Ethernet fabric, collectively called Data Center Bridging.

For additional information about how each of the mechanisms works, see “Ethernet Enhancements Supporting I/O Consolidation” at http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns783/white_paper_c11-462422.html.

The most basic way for FCoE to make Ethernet topologies lossless is to enable congested ports to send PAUSE control frames, which are specified in IEEE 802.3x. This technique instructs the transmitting port to stop sending data temporarily to avoid the need to drop packets. Using the PAUSE frame like this is a simple way to make the transmission lossless, but it can cause a ripple effect of congestion over the network, which is an impediment to performance as well as scalability. DCB extends the notion of quality of service (QoS) beyond a point-to-point scenario, covering the entire data center cloud.

Further, it provides a level of granularity that allows prioritization of specific traffic types. The primary technologies that enable DCB are:

- **Priority Flow Control** allows definition of up to eight user priorities on a single physical link, each of which has its own logical lane that can be paused independently of the others. This capability allows a port to send a Pause command with less effect on the network as a whole, because it enables more granular control over which traffic is affected, including application of lossless transmission only to FCoE traffic.
- **Congestion Notification** responds to congestion by communicating with upstream network devices, instructing them to throttle specific transmissions to shape the traffic that is causing congestion. That shaping is calculated to push the congestion to the edge of the network and away from the network core, limiting the effect of the congestion on network performance overall.
- **Enhanced Transmission Selection** allocates bandwidth among different traffic classes to help ensure appropriate priority and latency for traffic that requires it (for example, to help ensure that

bandwidth requirements are met for storage traffic to help ensure lossless transmission). The mechanism is flexible enough to allow other traffic classes to use idle allocated bandwidth, helping ensure overall efficient use of network resources.

- **Data Center Bridging Capability Exchange Protocol (DCBCXP)** is a management protocol that allows Enhanced Ethernet to work transparently with conventional Ethernet by dynamically discovering the capabilities of peer devices on the network. For example, DCBCXP enables an edge switch to discover the extent of the Enhanced Ethernet cloud and the related capabilities of its peers, so that it knows how to interact with them. DCBCXP also allows devices to verify that configuration parameters such as user priorities are compatible among devices and to push those parameters out to peers as needed.

These capabilities enable a robust Enhanced Ethernet topology that operates in parallel with traditional Ethernet over the same fabric for maximum capability and flexibility. Some of the capabilities provided by a unified data center fabric are listed in Table 2.

Table 2. Unified data center fabric capabilities.

Fibre Channel over Ethernet (FCoE)	Facilitates consolidation of Local Area Network (LAN) and Storage Area Network (SAN) traffic over a Unified Data Center Fabric
Data Center Bridging (DCB)	Enables Ethernet fabric to support lossless traffic flows, making it suitable for SAN traffic
SAN connectivity to all servers	FCoE enables all servers to access new or existing Fibre Channel SANs over standard 10GbE adapters, eliminating the need for Fibre Channel Host Bus Adapters and secondary storage fabrics
Enhanced virtualization infrastructure	Enables virtual machines to participate in live migration for load balancing and automatic failover, as well as network-interface virtualization for enhanced traffic segmentation
Robust ecosystem of development partnerships (including Cisco, Emulex, Qlogic, Intel, and others)	Delivers integration of capabilities throughout the Unified Data Center Fabric hardware/software solution stack
I/O consolidation and data center network convergence	Reduces power and cooling requirements, for lower operating costs and environmentally sound operation, in addition to lower infrastructure costs

Deployment Considerations for Unified Data Center Fabric

As discussed here, a hallmark of FCoE design (and a core strategy to help ensure its adoption) is its capability to be deployed incrementally within existing environments, minimizing disruption of the deploying organization while helping ensure maximum benefit. For this reason, FCoE is engineered specifically to coexist with existing topologies, such as traditional Ethernet, Fibre Channel, and iSCSI. To deliver additional flexibility, FCoE provides two distinct deployment scenarios, which implementers can combine, if needed, to ease the transition.

Flexible Deployment Options

The first of these deployment scenarios is the general case where both SAN and LAN I/O are passed over a single network interface card (NIC), as shown in Figure 4. Here, a single 10 Gigabit Ethernet network supports both types of traffic, simplifying the environment and enabling cost efficiencies.




Figure 4. FCoE enables unified Ethernet networking, passing SAN and LAN traffic over a single 10 gigabit ethernet NIC.




Figure 5. FCoE supports the passing of SAN and LAN traffic over separate 10 Gigabit Ethernet NICs.

As an alternative to the typical single-NIC model of FCoE deployment, servers can pass LAN and SAN traffic over separate 10 Gigabit Ethernet NICs, as shown in Figure 5. This scenario is appropriate for some organizations that want to obtain many of the advantages of the FCoE topology, but who need to physically separate the two networks, for example for organizational or regulatory reasons.

This second deployment option enables organizations to segment LAN and SAN traffic at the physical layer, as opposed to using virtual LANs (VLANs) or the user priority groups that are enabled by FCoE. This approach can have positive implications for organizations that rely on such physical segmentation to control network traffic or load balance the network, as well as for industries such as healthcare and financial services, where it may be used to achieve regulatory compliance.

FCoE-Capable Data Center Switches from Cisco

For either of these two deployment models, the Cisco Nexus 5000 Series offers feature-rich Layer 2 data center switches with advanced support for unified data center fabric. In addition to FCoE support over Enhanced 10 Gigabit Ethernet, the Cisco Nexus 5000 Series delivers advanced lossless transmission capabilities, taking advantage of the Enhanced Ethernet capabilities mentioned previously, helping ensure the fidelity of converged traffic. The switches also have a cut-through design that can deliver adapter-to-switch-to-adapter latency of less than 10 microseconds and port-to-port latency of approximately 3 microseconds, independent of packet size. The switches include ports at the rear for consistency with data center servers, allowing shorter and simpler cable runs within racks and reducing cost and copper waste. Front-to-back cooling is also consistent with server cooling designs, helping facilitate rack deployment and increase data center cooling efficiency.

10 Gigabit Ethernet Connectivity from Intel

Intel® 10 Gigabit Server Adapters for PCI Express® are designed to meet the throughput requirements of bandwidth-hungry implementations, such as those associated with unified data center fabric. Intel is delivering FCoE initiator support for Red Hat Enterprise Linux® and Microsoft Windows® Server, playing a pioneering role in bringing 10 Gigabit Ethernet connectivity for FCoE to the data center. The design of these adapters also incorporates DCB capability to provide the lossless, no-drop behavior required by FCoE. These adapters are optimized

The Cisco Nexus 5000 Series offers these features:

- Cisco Data Center Ethernet (DCE) feature suite for lossless congestion management
- FCoE support for I/O consolidation
- Flexible connectivity, including 10 Gigabit Ethernet and Fibre Channel, for simplified in-rack cabling
- Redundant, hot-swappable power and cooling systems for high availability
- Flexible management, including Cisco Data Center Network Manager (DCNM) (after first customer shipment [FCS]), Cisco Fabric Manager, a command-line interface (CLI), Simple Network Management Protocol (SNMP) support, and Extensible Markup Language (XML) support, for high compatibility
- Rear-facing ports and front-to-back cooling for consistency with server designs

for multi-core processors and virtualized environments, with support for multiple queues that distribute I/O processing among multiple processor cores and help alleviate I/O bottlenecks between virtual machines. Receive-side scaling (for Windows) and scalable I/O (for Linux) intelligently use multiple queues to increase the efficiency of direction of packet flows to specific processor cores for handling. Virtual Machine Device queue (VMDq) technology³ offloads data sorting and data copying from the Virtual Machine Monitor software layer to the hardware, improving overall throughput and CPU utilization on virtualized servers. Extended Message-Signaled Interrupt (MSI-X) technology passes interrupts from multiple queues to particular processor cores simultaneously, enabling better load balancing across cores for better CPU utilization and lower latency.

Intel® 10 Gigabit Server Adapters offer these features:

- Support for Enhanced Ethernet features such as Priority Flow Control that provide no-drop behavior on Ethernet
- Load balancing across CPUs for increased performance on multiprocessor and multi-core systems
- Virtual machine device queues (VMDq) for efficient routing of packets in a virtualized environment
- Intel® I/O Acceleration Technology for more efficient data movement through the processor⁴
- Capability to toggle between interrupt aggregation and nonaggregation modes for low latency with diverse data
- Network packet handling without waiting or buffer overflow for efficient packet prioritization
- Support for eight-lane connectivity using standard and low-profile PCI Express® slots for fast, efficient data transfer
- Single- and dual-port configurations for copper and fiber networks

Conclusion

A convergence of factors continues to increase the use of SANs and the adoption of 10 Gigabit Ethernet. These factors include the requirements of server virtualization technology, the rapid adoption of multisoocket servers hosting multi-core processors, and the ever-increasing demands of today's business applications. FCoE has arrived as an enabling technology that controls the cost and complexity of deploying SAN connectivity to all the servers in the business environment. Emerging standards-based technologies provide lossless data transmission over a converged Ethernet fabric that combines LAN and storage traffic over the same wire, reducing infrastructure and operating costs in the data center as well as helping protect the environment.

FCoE has been designed from the ground up for maximum flexibility, including the capability to coexist with existing infrastructure, management topologies, and networking technologies. Backing from industry leaders like Cisco and Intel helps ensure the continuing innovation that is pushing FCoE into the networking environments of companies of all sizes and across all industries. Innovative hardware advances like the Cisco Nexus 5000 Series and Intel 10 Gigabit Server Adapters put these technologies into practice, building future data centers that are increasingly capable, flexible, efficient, and environmentally sound.

THE NEXT STEP

Switching support from Cisco for standards-based FCoE, including the Cisco Nexus 5000 Series, brings this technology to market with the assurance that comes from an industry leader and a system of collaborators that will help make FCoE technology prevalent in data centers worldwide.

For more information about FCoE products and technologies from Cisco, visit www.cisco.com/go/nexus5000.

Innovative 10 Gigabit Ethernet LAN adapters from Intel deliver lossless Ethernet for FCoE and interoperability with products from various switch vendors. Intel leads the industry in delivering FCoE initiators for both Windows and Linux. Top manufacturers offer Intel® adapters in their new servers.

For additional product information about Intel® networking connectivity products, visit www.intel.com/network.

¹ Source: Internal Intel estimates.

² Computerworld, "Doing the Math," April 3, 2006.

³ Intel's Virtual Machine Device queue (VMDq) technology requires an operating system that supports VMDq.

⁴ Intel® I/O Acceleration Technology (Intel® I/OAT) requires an operating system that supports Intel I/OAT.



Intel Corporation
2200 Mission College Blvd
Santa Clara, CA 95052
USA
www.intel.com

Copyright © 2008 Intel Corporation. All rights reserved.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site at <http://www.intel.com/>.

Intel, the Intel logo, Intel. Leap ahead., Intel. Leap ahead. logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.
0408/BY/MESH/PDF Please Recycle 319745-001US



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, CCENT, Cisco Eos, Cisco Lumin, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn is a service mark; 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, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, 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. (0804R)