Lossless 10 Gigabit Ethernet: The Unifying Infrastructure for SAN and LAN Consolidation

Introduction

As organizations increasingly rely on IT to help enable, and even change, their business strategies, they need their IT infrastructure to be more powerful, agile, and cost effective than ever. Today's enterprises require continual system availability, demand ubiquitous access, and expect rapid and fluid responses to their ever-changing business needs.

To provide these functions, enterprise data centers are challenged to get more use out of existing resources and operate quickly, with increased agility. Specifically, they must address the following data center challenges:

- Improve asset utilization to reduce or defer capital expenses
- Reduce power and cooling consumption to cut costs and align with green business practices
- Reduce time to market for infrastructure and applications
- Make data and resources available in real time to provide flexibility and alignment with current and future business needs

To meet these challenges, organizations must build a single network infrastructure that unifies traditional server I/O, storage, and network operations to more efficiently support evolving business applications. This document describes how a Unified Fabric, proposed by Cisco and NetApp, enables convergence of SAN traffic over an Ethernet and compares the performance of Fibre Channel over Ethernet (FCoE) to native Fibre Channel.

Today's Storage Network

Most large data centers have a separate, dedicated storage network, called a SAN, using Fibre Channel. Fibre Channel SANs connect servers to shared block storage systems through a dedicated high-speed network of host bus adapters (HBAs) and switches. The resulting storage network allows servers and storage to communicate, and multiple pathways can be established to help ensure storage availability by enhancing redundancy and improving performance. SANs have traditionally employed the Fibre Channel protocol using the physical implementation and signaling detailed in ANSI standard X3.230-1994 (ISO 14165-1). Fibre Channel uses optical fiber, coaxial copper, or twisted-pair copper cabling to carry SAN data over interswitch links (ISLs) at speeds of 1, 2, 4, and 10 Gigabits (Gb), and more recently, 8 and 20 Gb. Fibre Channel can operate in point-to-point, switched, and loop modes, with switched mode prevalent.

The cost and complexity of Fibre Channel kept SAN deployments out of reach for most small and midsized businesses (SMBs) until the introduction of SANs based on the Small Computer System over IP (iSCSI) protocol ratified by the IETF in 2003. An iSCSI SAN typically employs network interface cards (NICs), an iSCSI software driver, and Ethernet switches in a segregated network. Every major operating system includes an iSCSI driver in its distribution. iSCSI performance can be improved by deploying Ethernet NICs with a TCP/IP offload engine to reduce the CPU demands for TCP/IP processing. Given that iSCSI uses the standard SCSI command set and IP network, interoperability is well understood and straightforward. Currently, the vast majority of iSCSI SANs operate at 1 Gb Ethernet speeds; however, 10 Gigabit Ethernet NICs, switches, and storage systems are available today, and as costs continue to come down, 10 Gigabit Ethernet solutions will become more widely deployed.

A third popular form of storage is network attached storage (NAS). NAS consists of a server connecting the storage to the TCP/IP network and enabling users to access the storage by using either the Microsoft Common Internet File

System (CIFS) or UNIX Network File System (NFS) protocol over TCP/IP. Typical NAS applications include file serving, file sharing, home directories, technical applications such as software development and electronic design automation, database storage and grid computing. An iSCSI SAN provides block-level access to data (it uses a disk drive), whereas NAS provides file-level access to data (it serves up a file).

Requirements for Unified Infrastructure

Today, most organizations have to run parallel network infrastructures for their LANs and SANs, with separate switches and separate host bus adapters. For Fibre Channel SANs, one or more 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, increasing costs even more. With Fibre Channel SANs, separate networks must be operated for the LAN and SAN environments, as shown in Figure 1.

Figure 1. Today Most Data Centers Use Separate LAN and SAN Networks with Separate Switches and Network Adapters in the Server; This Topology Requires a Dedicated Fibre Channel Infrastructure

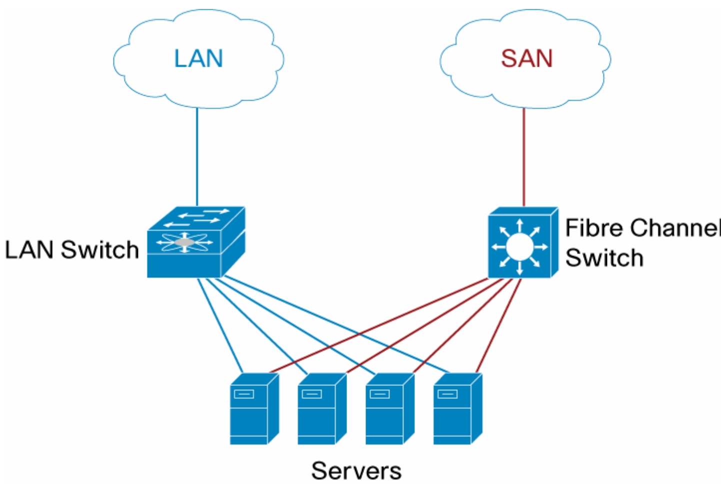

These separate networks require added expense due to requirements such as increased number of network interfaces, additional cabling and switch ports, and more complex support needs. Another factor in increase number of network adapters is server virtualization. Server virtualization, such as that provided by VMware, requires multiple adapters to carry traffic for LAN, SAN, hypervisor management, and virtual infrastructure services.

As the environment grows over time those expenses become even greater. Currently, the server environment and the server access layer of the network are particular areas of focus. Because of the scale of the server environment, with hundreds or even thousands of servers, small changes can have significant effects.

Lossless 10 Gigabit Ethernet: The Unifying Infrastructure

The storage network can be classified into two categories: storage for IP-based traffic such as iSCSI, NAS, Internet Fibre Control Protocol (iFCP), and Fibre Channel over IP (FCIP); and Fibre Channel based storage. Storage for IP-based traffic uses TCP/IP as the underlying protocol and therefore could be deployed over the same infrastructure as the LAN. Fibre Channel–based storage requires a separate network infrastructure which increases capital expenditures (CapEx) and operating expenses (OpEx) and leads to inefficient utilization of network resources.

Ideally, both LAN and SAN traffic would be supported by a single network fabric that would support the coexistence of different storage network models without compromising their operations or management. This would enable IT managers to optimize flexibility and efficiency, while minimizing any risk and disruption to their existing data center environment. Fortunately with the advent of FCoE, infrastructure unification is possible, as Ethernet becomes the common transport for the Unified Fabric. FCoE is a new technology and emerging standard that allows transmission of Fibre Channel traffic over Ethernet fabric by encapsulating native Fibre Channel frames into Ethernet packets, as shown in Figure 2.

Figure 2. In FCoE Packets, the Fibre Channel Payload Is Encapsulated in the Ethernet Header

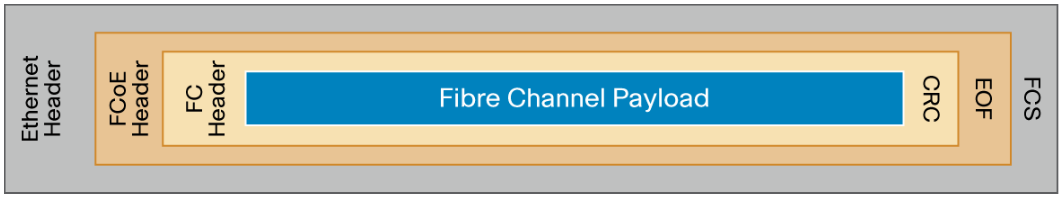

This approach preserves the native format of Fibre Channel frames, and FCoE traffic appears as Fibre Channel traffic to the Fibre Channel fabric. This characteristic allows IT departments to maintain the existing environment's operations, security, and traffic-management attributes before and after migration; it also preserves investments in Fibre Channel expertise and equipment.

FCoE enables the consolidation of the LAN and SAN over a common infrastructure that reduces the number of server I/O adapters and cables, lowers power and cooling costs through the consolidation of server I/O adapters, and eliminates unnecessary switching infrastructure. With the capability to transport Fibre Channel over Ethernet, the common infrastructure for storage is in place: 10 Gigabit Ethernet. Although 10Gbps Ethernet provides a solid foundation for a Unified Fabric, there is one additional characteristic that it must provide to enable the transport of Fibre Channel over Ethernet and that is lossless operation.

Why lossless 10 Gigabit Ethernet? Lossless Ethernet is a prerequisite for consolidation of Fibre Channel SAN traffic over a Unified Fabric. A set of network technologies, including IEEE Data Center Bridging, enables Ethernet fabrics to support lossless transmission, making them suitable for carrying all types of SAN traffic. This set of emerging standards enables better traffic prioritization over a single Ethernet interface and provides an advanced means for shaping traffic on the network to decrease congestion. A description of how Cisco Unified Fabric and IEEE Data Center Bridging enables convergence of different I/O traffic onto a single Unified Fabric in the data centers is presented in the next section.

Cisco Nexus 5000 Series Switches

Cisco Nexus™ 5000 Series Switches (Figure 3) are the devices in the Cisco® data center switch portfolio that deliver on the promise of a Unified Fabric for the data center. This Unified Fabric will have the operational characteristics to concurrently handle LAN, SAN, and server clustering traffic. In addition, the Cisco Nexus 5000 Series has a cut-through design architecture that can deliver a consistent port-to-port latency of 3.2 microseconds independent of the packet size. The Cisco Nexus 5000 Series' high performance line rate 10 Gigabit Ethernet throughput characteristic combined with its low latency enable a outstanding efficiency and performance for storage networks. Low latency and low jitter are essential requirements for the high-performance computing applications that the Cisco Nexus 5000 Series consolidates over the same Unified Fabric.

Figure 3. Cisco Nexus 5020 Unified Fabric Switch

Consolidating the multiple networks in the data center into a single network will promote significant reductions in capital and operating expenses. Just as important, giving every server in the data center consistent and ubiquitous I/O capabilities will markedly improve the service capabilities, flexibility, and resilience of the data center.

One of the primary enablers of Unified Fabric is a set of enhancements to Ethernet, defined within the IEEE Data Center Bridging standards body, which the Cisco Nexus 5000 Series supports. Cisco Unified fabric is an architecture based on a collection of open standard Ethernet extensions to improve and expand Ethernet networking and management capabilities in the data center. These standards provide congestion management capabilities that enable Ethernet to be a lossless transport, thereby facilitating I/O consolidation on a Unified Fabric.

Currently, the distinct interconnect technologies transport different traffic from different applications. The primary Ethernet extensions are as follows:

- Priority Flow Control (PFC) 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 frame for a given class of service that requires a lossless transmission without affecting other traffic. For example FCoE traffic uses PFC to create a lossless transmission path. This capability enables lossless 10 Gigabit Ethernet.
- Enhanced Transmission Selection (ETS) 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). 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 Exchange (DCBX) is a management protocol that allows Data Center Bridging capable switches to work transparently with conventional Ethernet switches by dynamically discovering the capabilities of peer devices on the network. For example, DCBX enables an edge switch to discover the related capabilities of its peers, so that it knows how to interact with them. DCBX also allows devices to verify that configuration parameters such as user priorities are compatible among devices and to send those parameters out to peers as needed.

The unified network fabric delivered with the Cisco Nexus 5000 Series will help organizations align their IT assets with their business priorities and deliver tangible business benefits. Table 1 summarizes these benefits.

Table 1. Cisco Nexus 5000 Series Benefits

Benefit	Description
Lower total cost of ownership (TCO)	The Cisco Nexus 5000 Series provides a Unified Fabric over Ethernet for LAN, storage, and server cluster traffic. This Unified Fabric provides consolidation and higher utilization of previously separate resources, reducing the number of server I/O adapters and cables needed by up to 50% and lowering power and cooling costs significantly through the elimination of unnecessary switching infrastructure.
Investment protection	The evolution to a Unified Fabric will happen gradually as compelling business and technical factors present themselves. The Unified Fabric provided by Cisco Nexus 5000 Series Switches through FCoE allows organizations to preserve the investment they already have in their existing Fibre Channel infrastructure during this transitional period. The Cisco Nexus 5000 Series is designed to allow customers to take advantage of the immediate benefits of a Unified Fabric without compromising architectural, management, or operational best practices in either Ethernet or Fibre Channel networks.
Increased business agility	The Cisco Nexus 5000 Series supports Cisco VN-Link, enabling applications to move across the network while maintaining their provisioned network services, such as security policy, quality of service (QoS), and overall performance. As a result, IT managers can easily move applications based on factors such as CPU utilization and thermal loads while still preserving their network characteristics. This capability allows IT departments to rapidly respond to changing business demands through rapid provisioning of the dynamic and flexible application infrastructure enabled by the combination of server virtualization and the Cisco Nexus 5000 Series virtual machine–optimized services.
Enhanced business resilience	The Cisco Nexus 5000 Series is engineered with specific hardware and operating systems to deliver component- and system-level operational continuity.

NetApp Unified Storage Solutions

NetApp FAS systems function as storage unification engines within the data center that can simultaneously support multiple network protocols, including Fibre Channel, FCoE, iSCSI and NAS. IT departments consolidate and simplify their storage environments using NetApp technology. Now, this unified storage system design from NetApp coupled together with the Cisco Unified Fabric design allows the deployment of a single, integrated storage and network solution that supports a broad range of server connections (Figure 4).

Figure 4. NetApp FAS System Unifying Fibre Channel, IP SAN, and NAS Storage

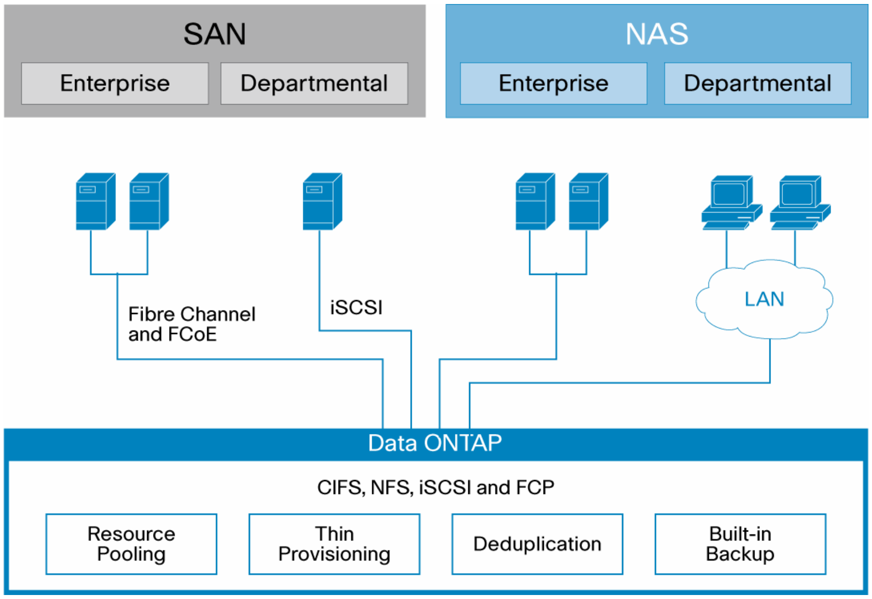

Performance Analysis of FCoE with Cisco Nexus 5000 Series

With the support for native FCoE in both the servers (FCoE initiators) and the NetApp storage system (FCoE target), the Unified Fabric enabled by the Cisco Nexus 5000 Series provides the capability to consolidate the SAN and LAN without risking any negative effect on the storage environment. The capability of the Unified Fabric to provide lossless behavior and guaranteed capacity for the storage traffic helps ensure that the storage I/O is protected and has the necessary capacity and low latencies to meet the critical data center requirements.

In a test comparison of the two environments (Figure 5), one a traditional distinct LAN and Fibre Channel SAN configuration and the other using the FCoE capabilities of the Unified Fabric, the servers' access to the storage environment was identical in response times, throughput, and I/O operations per second (IOPS).

Figure 5. Test Topology Comparing Native FCoE Through the Cisco Nexus 5000 Series with Native Fibre Channel

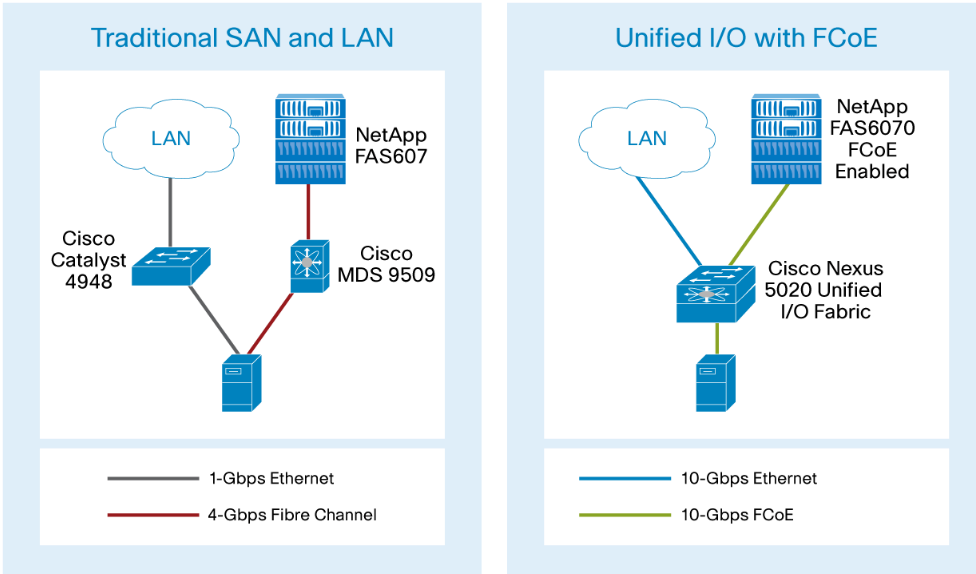

As shown in Figure 6, the FCoE environment has the same throughput performance for both read and write operations as the Fibre Channel environment.

Figure 6. Test Results Indicating That FCoE Has the Same Throughput Performance for Both Read and Write Operations as Native Fibre Channel

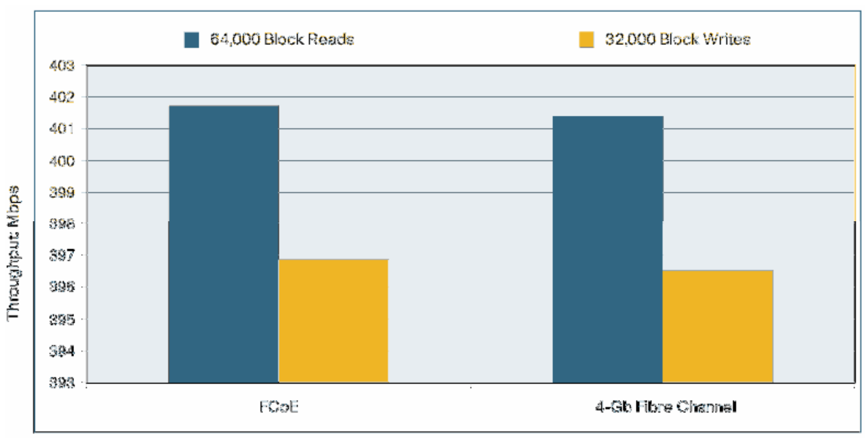

The results demonstrate that the end-to-end behavior of the storage system I/O from the server and application perspective is identical in the consolidated unified I/O environment to that of the standalone Fibre Channel configuration. This is understandable given that the endpoints, the server and filer, use identical block read and write operations in both cases, and the only difference is that with FCoE, the Fibre Channel communication is being carried across the Unified Fabric.

The current generation of converged network adapters (CNAs) for FCoE support up to 4 Gbps of Fibre Channel traffic. The HBA driver and hardware on the shipping CNA is identical to that on the 4-Gbps Fibre Channel HBAs. The next generation of CNAs will support the 8-Gbps Fibre Channel standard.

Test Environment

The physical topology of the test environment is Figures 5 and 6. To provide an equal comparison, the Fibre Channel network was limited to a single hop (server and filer connected to the same switch).

Storage traffic was generated and measured using the NetApp Simulated IO (SIO) workload generator: see http://partners.netapp.com/go/techontap/tot-march2006/0306tot_monthlytoolSIO.html.

Table 2 summarizes the test environment.

Table 2. Test Environment

	Hardware Version	Software Version
Unified I/O switch	Cisco Nexus 5020 Switch	4.0(0)N1(1a)
Storage System	NetApp 6070	ONTAP 7.3.1
Gigabit Ethernet Switch	Cisco Catalyst® 4948 Switch	12.2(46)SG
Fibre Channel Switch	Cisco MDS 9513 Multilayer Director	3.2(2c)
Server	HP DL585, 4X Opteron CPU, 32 GB DRAM	LINUX
can	QLogic QLE8042	4.3.1

Deployment Model for Unified Fabric Using Cisco Nexus 5000 Series and NetApp Unified Storage Solutions

The transformation to a unified network can happen incrementally within existing environments, reducing disruption for the organization while helping ensure optimal benefit. In the initial phase, the access layer is the area of the focus, with the Cisco Nexus 5000 Series unified access layer switch providing a unified fabric for LAN traffic, IP-based storage traffic, and Fibre Channel–based storage traffic. 10 Gigabit Ethernet will be the natural common infrastructure for Unified Fabric.

In the deployment scenario shown in Figure 7, the Cisco Nexus 5000 Series is used as the Unified Fabric, consolidating the LAN and SAN at the access layer, and NetApp FAS is used as the storage unification engine, simultaneously supporting the Fibre Channel and the IP SAN and NAS. The unified storage and fabric infrastructure reduces acquisition and operational costs by reducing the number of point solutions needed.

Figure 7. Unified Network Solution with Cisco and NetApp; the Cisco Nexus 5000 Series Provides Unified Fabric for the LAN and SAN at the Access Layer, and NetApp FAS Provides Unified Storage

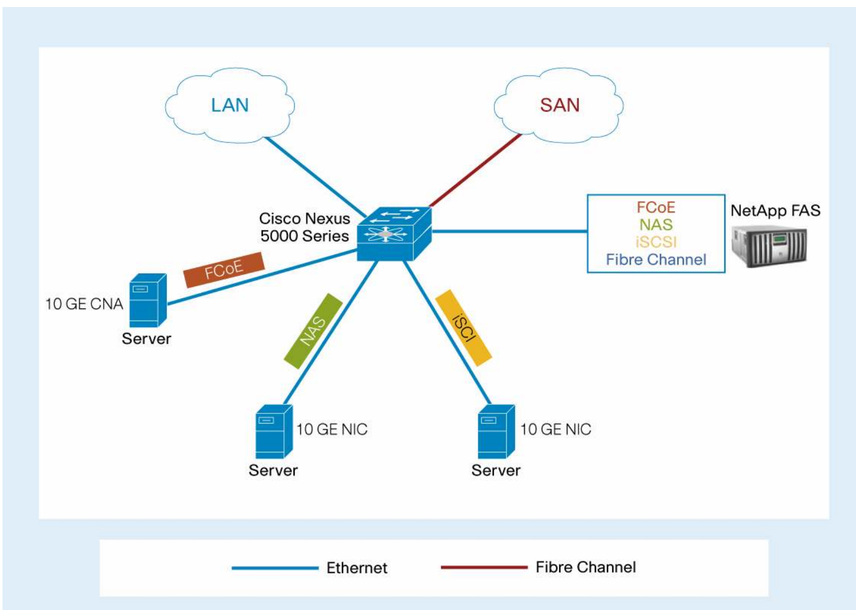

The deployment scenario shown in Figure 8 uses the Cisco Nexus 5000 Series for the Unified Fabric at the access layer to consolidate the access layer switch, where there are two parallel LAN and Fibre Channel SAN infrastructures. From the servers, both SAN and LAN I/O are passed over a single 10 Gigabit NIC called a converged network adapter, or CNA. This converged network reduces the number of adapters needed, greatly simplifies cabling, and reduces the number of access layer switches needed, enabling cost efficiencies.

Figure 8. Today Data Centers Use Separate LAN and SAN Networks with Separate Switches and Network Adapters in the Servers; the Unified Fabric Data Center Uses a Consolidated LAN and SAN Network at the Access Layer with a CNA in the Servers and the Cisco Nexus 5000 Series as the Unified Fabric Switch

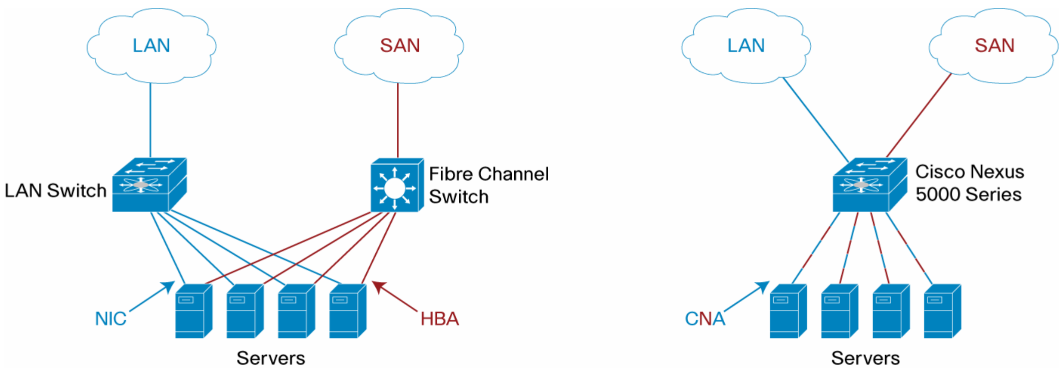

Conclusion

Today data centers deploy multiple networks, based on distinct interconnect technologies, to transport traffic from different networks such as LANs, SANs, and server-to-server interprocess communication (IPC) systems. A typical server in an enterprise data center therefore has multiple interfaces to allow it to be connected to the various disparate networks. With the nearly universal mandate across all IT departments to reduce operational complexity and cost, managing multiple interconnect technologies to meet data transport needs for different applications is no longer an optimal business and technology practice. With recent advances in speed over Ethernet (10-Gbps Ethernet is already in deployment, and 40- and 100-Gbps Ethernet are in development), Ethernet has become an attractive choice as the technology for convergence in the data center. Moreover, new enhancements in Ethernet and the advent of FCoE enable consolidation of different storage solutions and traditional separate networks over the same infrastructure. Tests show that the performance of FCoE is equivalent to that of native Fibre Channel with the Cisco Nexus 5000 Series used to access NetApp unified storage.

Cisco and NetApp together enable customers to deploy a Unified Fabric solution. NetApp Ethernet-based storage systems, through unified, multiprotocol architecture, provide storage consolidation, and the Cisco Nexus 5000 Series Unified Fabric, based on emerging standards technologies, provides lossless data transmission over a converged Ethernet fabric that combines LAN, storage, and high-performance low-latency traffic on the same infrastructure.

About NetApp

NetApp creates innovative storage and data management solutions that accelerate business breakthroughs and deliver outstanding cost efficiency. Discover NetApp's passion for helping companies around the world go further, faster at http://www.netapp.com.

For Further Information

For more information about the Cisco Nexus 5000 Series and unified fabric solution, visit http://www.cisco.com/go/nexus5000.

For more information about the NetApp unified storage solution, visit http://www.netapp.com/us/products/protocols.

Printed in USA C11-501770-01 05/09