Storage Networking 101

Abstract

This document describes storage networking methods, including history, concepts, and designs, from a high-level introductory stance. The primary objective of this paper is to provide the reader with a baseline understanding of storage networking as preparation for more advanced storage concepts.

Introduction: Storage and Networks—In the Beginning

Since the first computer was developed, there has been an increasing demand to make them faster, cheaper, and more applicable to everyday lives. The early mainframe computers have evolved from large, centralized systems to more nimble, enterprise-class servers brought on by cheaper and more efficient computing technologies. In turn, networking technologies have had an effect on the evolution of computing platforms. The maturation of two technologies and the increasing hunger for computer processing power and associated data have driven the need for faster, more accessible data storage techniques and the advent of storage networking.

Historically, storage devices have always resided inside or behind servers and were tightly coupled with these servers' operating systems. Connection between servers and storage devices were typically facilitated using the (Small Computer Systems Interface (SCSI) protocol. SCSI is a storage interface that has become the dominant standard for storage devices such as disk and tape drives inside and outside of high-performance servers. Because of its high data transfer rates, reliability, and low latency SCSI is an ideal protocol for connecting various storage devices to a server.

In the past 10 to 15 years, changes in the ways business is conducted and the explosion of Internet-based business practices has resulted in additional challenges that affect the way information is accessed and stored. The resulting increased need for storage capacity has placed heavy pressure on IT organizations to keep up with the demand for storage, while expanding its capacity in a cost-effective and manageable fashion.

Introduction to Computer Storage

Computer storage holds the lifeblood of today's Internet-based economy: *information*. Information and the knowledge derived from it are the core elements by which our society increases its productivity and grows at unprecedented rates. Along with the explosion of the Internet and the creation of our information-based economy, digital data needs to be stored and made available to those systems needing access to them. Computer storage comprises various technologies and media that enable information storage, access, and protection.

The most common types of information storage use magnetic media to save and retrieve imbedded data. Magnetic media is made up of two primary formats—fixed and removable. Fixed is commonly associated with disk storage technology where the read/write heads and the spinning media (disk platters) are held together in one unit. Removable media is commonly associated with tape technology where the read/write heads are contained in a "transport" (or drive) and the cartridge (or tape reels in older systems) hold the tape that is sequentially passed across the heads to copy or retrieve the data. The inexpensive tape media can also be removed and stored either onsite or offsite locations for improved data protection.

The primary benefits of magnetic media are its flexibility (fixed or removable), relatively low cost, good performance, and its ability to be reused easily. This media can be "rewritten" with different digital information and safely stored for many years and accessed in the future. One drawback of magnetic media is that it's a physical system with moving parts that can wear and fail. Also, natural disasters and system failures can destroy data or make it unavailable to those applications needing access it. This inherent risk places a large responsibility on the storage system manager to implement ways to protect their valuable information assets. The traditional method of protecting data utilizes various forms of replication and parity error protection. The idea behind this is that the more redundant the data is the more likely you are to be able to retrieve it before it becomes lost or corrupted. Each media format has its strengths and weaknesses in the areas of price/performance, scalability and reliability.

The primary magnetic storage used for online, server-accessible storage is disk. Disk storage comes in two primary configurations: JBOD and RAID. JBOD stands for Just a Bunch of Disks and represents the simplest and least expensive "raw storage" option. The individual disks are arranged in a simple cabinet and are available to its servers as a group of independently accessible disks. They have little or no "buffering" (or cache memory) or an intelligent controller that enables "striping" (or replicating data across disks) for improved performance or data protection. JBOD drives share power and the physical cabinet with one another. JBOD has limited growth capacity and usually scales to much less than 1 terabyte (TB) per cabinet. Since they have no inherent intelligence, parity checking or data striping, there is no protection in the event of a drive failure. Tape storage would be a common method of saving information in the event of a drive failure.

In contrast to JBOD, RAID storage stands for Redundant Array of Independent (or Inexpensive) Disks and is controlled by an intelligent controller (usually with large amounts of solid state memory) that provides parity checking and enables data "striping" (or replicating the same data) across multiple drives for improved protection and performance. In addition to parity checking and data striping across drives, critical business data can be protected by "mirroring" specific data volumes within an array cabinet or across multiple arrays. Disk mirroring provides an identical copy of the data in the event of a failure occurring on the production data volume. This fault-tolerant grouping of disks can present itself as a single disk volume to the servers attached to it. RAID storage commonly scales from a few drives per cabinet up to 512 drives per cabinet. That translates to 10's of terabytes per cabinet depending on the manufacturer, density and size of drives used, and levels of RAID protection desired. The basic unit of storage in each of these configurations is the HDAs (or drives) that are inserted into each cabinet bus to enable it to read/write data. These drive capacities range from 18, 36, 72 and now up to 181 GB of data storage per drive with that capacity increasing rapidly. RAID storage delivers increased storage capacity, performance, and availability over simple JBOD storage configurations.

There are seven basic levels RAID data protection implemented in storage arrays (or subsystems). RAID levels 0 (zero) through 6 were defined in the original University of California Berkeley RAID project. RAID 2, 4 and 6 are rarely seen in commercial products. RAID 0 is merely disk striping and offers some performance advantages. Unfortunately, it doesn't store parity information thus does not offer true RAID protection. RAID 1 offers complete duplication of data and 100% data redundancy but it can be too expensive for most applications. RAID 3 and 5 each use one extra disk to store parity information needed to recreate data in the event of a single disk failure. RAID 3 uses a dedicated parity disk and is typically faster for throughput-oriented, sequential applications. RAID 5 distributes the parity information across all disks in the array and is typically faster for transaction oriented and random access applications. Adding cache memory to the controller (up to 32GB on some models) greatly improves performance for all RAID levels by making commonly accessed information in solid state memory verses reading/writing the data from rotational disk or sequential tape.

To protect customers' critical business information stored within large online disks (that is, fixed) storage, removable tape media is commonly used. Traditionally, application software is run on each server that copies the data (made up of files or volumes) from the disks and places it on tape. This magnetic tape is commonly held within a plastic cartridge with capacities ranging from 800 MB (IBM 3480 format), 20 GB (for the popular STK 9840), to 100 GB (for the new Ultrium and LTO tape formats). Most large companies implement automated tape libraries, which use high-speed robotics (using arms and fingers) to retrieve specific cartridges stored in holding cells and place them into the tape transport (or drive) to be written or read. After the read/write process is completed, the tape is ejected and placed back into the cell or may be removed from the library and moved to an offsite location for data protection and archiving purposes.

Customers must combine both disk and tape technologies together to address their exploding requirements for information storage, access and protection.

Storage Networking Fundamentals

SCSI

The SCSI is the most common methodology for connecting storage devices to servers. SCSI emerged in 1979 as an 8-bit parallel bus interface with support for one or two disks. This protocol has evolved over the years, increasing its scope as the foundation for other storage related technologies. Today, serial SCSI is a layered and well-architected suite of protocols for requesting services from storage devices.

The SCSI Architecture Model (SAM-2), which is part of the National Committee for Information Technology Standards (NCITS) T10 standard, provides a layered model for SCSI implementation. SCSI driver software, physical interconnections, command implementation and storage management provide the framework for SCSI interoperability and scalability. There is support for multiple device types, queuing, multi-tasking, improved performance, caching, cabling and termination, automatic device ID configuration and dual port operation. Device-specific commands are organized into block (hard disk, optical devices, and other direct-access devices), streaming (tapes and sequential access devices), and other SCSI device-type specific command sets (defines commands for multimedia, optical jukeboxes and controllers). SCSI-3 command set (see below) converts the logical layer into a packet-based format that can be sent over a network. There are multiple implementations of serial SCSI including Fibre Channel, Apple's Firewire or 1394, IBM Serial Storage Architecture (SSA) and most recently iSCSI.

SCSI Standards provide for three possible electrical configurations:

- Low-cost "single-ended" alternative intended for devices in close proximity to each other (single cabinet) up to 6 meters
- More expensive "differential" alternative called "High Voltage Differential" (HVD) supports distances up to 25 meters with better noise immunity
- Most recent introduction (1995) of Low Voltage Differential (LVD) for SCSI-3 provides distances of up to 12 meters

With the growth of Internet-based applications, the accelerating information requirements dictated that the storage capacity grow at a faster rate than the server's processing power. Because a server's cabinet has a limited internal storage expansion capacity, the server's storage needed to be "externalized" to enable it to scale to meet the needs of the applications running on it. SCSI protocol and cabling was used to connect external storage devices to their associated servers to address this need. As servers and storage grew, managing both environments was an increasing challenge. Consolidation of servers and storage helped improved the management model but physical SCSI cabling was posing limits to the distance (25 feet maximum), speed and sharing capability of storage devices. Although the SCSI protocol is popular and well supported by both servers and devices, its physical distance and sharing capability is severely limited. Other technologies (such as Fibre Channel, Ethernet, IP and iSCSI) came alongside SCSI to overcome these limitations to make storage more accessible and manageable within the customer's environment. Now customers can gain the widespread support for SCSI with the benefits of more robust networking topologies. In spite of the limitations of physical SCSI, the protocol will continue to flourish as a high performance storage interconnect. The SCSI protocol architecture is well designed to support newer technologies to meet the ever-increasing demands of the customer for storage.   It is this innate resourcefulness for adaptability to newer technologies that has kept the SCSI protocol at the forefront of storage technology for high-end, scalable and reliable data access.

TCP/IP

Transmission Control Protocol and Internet Protocol (TCP/IP) make up the communication protocol suite. These protocols were used to develop the Internet infrastructure that is so prevalent. The suite's roots were based largely around the U.S. government's need to interconnect various defense departmental computer networks. Developed in the early 1970s, it has been expanded upon and continues to grow as new technologies evolve.

TCP/IP has played a major role in the success of the Internet. Besides having the ability to scale in very large network environments, TCP/IP fosters the ability for subscribers to share information safely and reliably with other users. These characteristics provide a truly open network framework to support millions of individuals in homes, schools, governments, small businesses and corporations in the remote corners of the world. With its support for a wide-variety of network technologies, TCP/IP is fully capable of providing the underlying foundation for global storage networks.

Ethernet

Ethernet is the most widely used local area network (LAN) technology today. Specified in the IEEE 802.3 standards, it was originally conceived and developed by Xerox Corporation. Its identification as an ideal technology for attaching desktop computers to one another led to additional development efforts by Intel and Digital. Since its inception, initial speeds of 10 Mbps have evolved into Fast Ethernet or 100BASET, providing transmission speeds up to 100 megabits per second. Gigabit Ethernet provides an even higher level of backbone support at 1000 megabits per second (1 gigabit or 1 billion bits per second) primarily on fibre-optic cable. Now, 10 Gigabit Ethernet is just around the corner and waiting to meet the next challenge of performance and connectivity.

Initially, governmental agencies and later educational institutions embraced Ethernet for their communication needs. Ethernet has proven to be a reliable transport that scales well to support a wide variety of applications and operating systems.

The combination of 10 Gigabit Ethernet and optimized TCP/IP processing in hardware provides a very compelling solution for supporting storage networking applications.

Fibre Channel

Most storage-area networks (SANs) are predominantly based on a Fibre Channel (FC) architecture. FC was developed to address the speed, capacity, and reliability requirements associated with communication between storage and server devices. It operates at 1 Gbps today with 2 Gbps currently in the testing phase. It has a sustained throughput of 100 MB/sec half-duplex and 200 MB/sec full duplex.

Fibre Channel development started in 1988 and the NCITS T11: I/O interface (X3.230-1994) standard was completed in 1994. It offered a solution to IT professionals who required improved reliability, higher performance and increased scalability when building the storage infrastructure. Fibre Channel storage area networks (SANs) provided servers the access to storage using the storage protocols and server-to-server communications using IP network protocols.

SCSI is often mapped as a higher layer protocol on Fibre Channel, leveraging the benefits of this reliable channel-based protocol with Fibre Channel's ability to resolve its shortfalls in speed, number of devices, and physical distance limitations.

Figures 1 through 4 below show the various topologies in which Fibre Channel devices can be configured to build storage networks.

Figure 1 Fibre Channel Loop using ring configuration

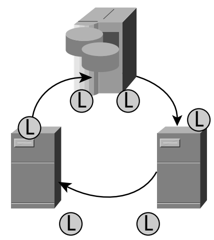

Figure 2 Fibre Channel Loop constructed using Fibre Channel hub

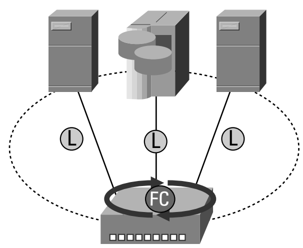

Figure 3 Switched Fibre Channel topology using node, fabric, and trunk ports

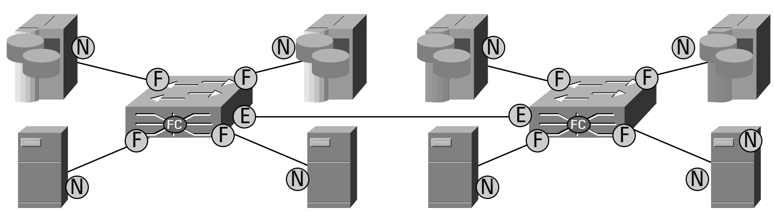

Architecture Basics

LAN/MAN/WAN

Figure 4 Graphical example of a network topology showing Local, Metropolitan, and Wide Area Networks. Note that this example assumes that MAN is providing access to WAN.

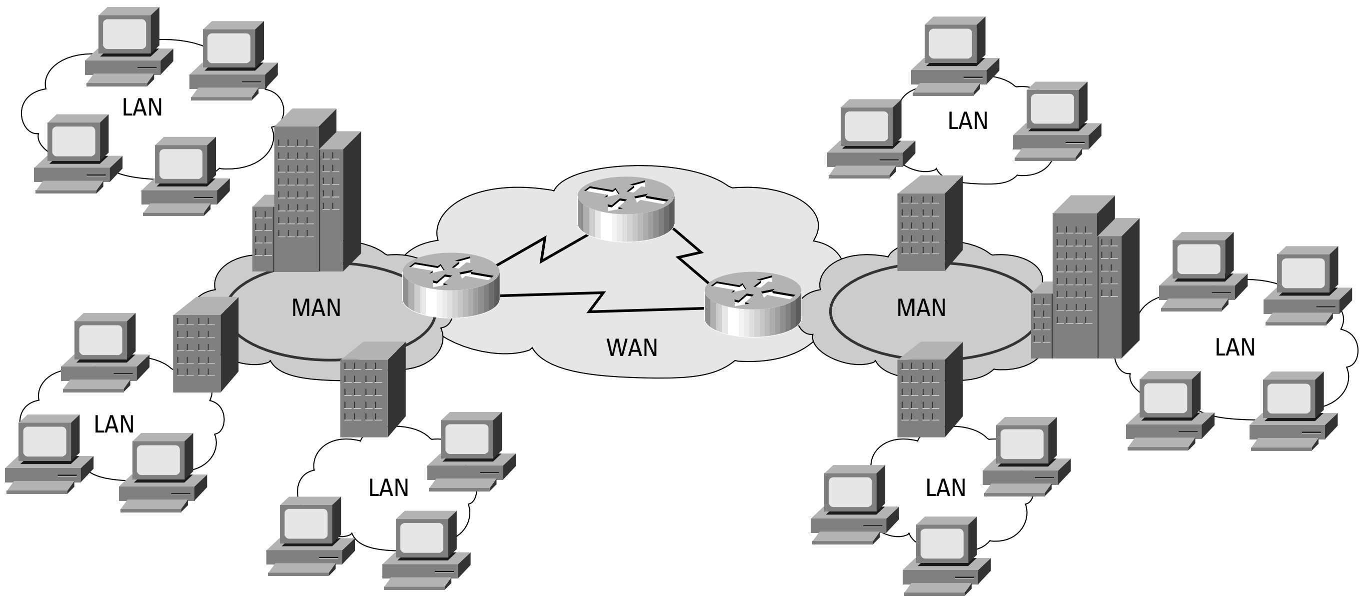

Local Area Network (LAN)

A LAN is a computer network used to connect personal computers, workstations, and printers over a relatively small area such as within an office or building. This network allows each attached computer to send/receive e-mail, share networked printers, and access applications running on file servers. The typical LAN architecture is built upon a desktop client and server model where network access is accomplished through logging into the server from the client. LAN equipment is typically connected together using copper twisted-pair wiring (Category 5), however short-range wireless connections are becoming popular. LANs typically run at speeds up to 100 Mbps. LANs are typically limited to the confines of a building or small campus, but are capable of being extended over MAN or WAN networks to remote users.

Metropolitan Area Network (MAN)

MANs are high-speed networks which are capable of spanning an area between 5 and 100 KM in diameter; the typical size of a major city. MANs are generally used to span distances outside the range of LANs and where there is not a need for the long haul distances characterized by WANs. Metropolitan network services and equipment are typically contracted from service providers who own the network, however it is becoming increasingly popular for larger companies to lease just the physical cabling (copper or fibre) in order to build their own private MANs. Metropolitan networks are capable of running at speeds up to 2048 Mbps and are typically architected using a point-to-point or ring topology. Advances in fibre-optic multiplexing of light wavelengths (dense wavelength-division multiplexing or DWDM) and the cost associated with leasing "dark" fibre is leading to a surge in the deployment of private MANs. MANs can be used to extend LANs or SANs across a city and also can provide a centralized connection to a WAN, as depicted in Figure 4 above.

Wide Area Network (WAN)

A WAN comprises a geographically dispersed grouping of network equipment owned and administered by a public carrier or network service provider. Leased, end-to-end, long-haul connections are provided to the customer with some guarantee of service levels. The primary factor determining the cost of the line is the distance spanned, the speed of the circuit and the level of competition. Common leased connections are either T-1 or T-3 lines, which provide speed in multiples of 64-kbps channels up to a maximum speed of 1.544 Mbps (24 multiplexed channels) or 45 Mbps (equivalent to 28 T-1 trunks) respectively. In Europe, these are described as E1 and E-3, both offering slightly more bandwidth than the domestic T series standards. Sometimes customers opt for higher speed services such as ATM OC3 at 155 Mbps or higher speed Synchronous Optical Networking (SONET) services. WANs are used when connecting resources over distances that exceed the reach of MANs. WAN topologies are normally arranged in either a star or mesh pattern. LANs or MANs are often connected over WAN in order to connect LANs or MANs in other locations or to provide access to such public resources such as the Internet.

Direct Attached Storage

Direct attached storage (DAS) was developed in the mainframe environment where networking was quite simple. In the 1980s, computing shifted from large centralized system into more flexible and networked client/server distributed model. Local area networking was in its early phases and accelerated this new shift in computing. Server-attached storage (SAS) was similar to DAS but with a distributed approach using the LAN as the connection. With increases in compute power, memory, storage density, and network bandwidth, more and more data was stored on PCs and workstations. Distributed computing and storage growth have proliferated and are driving the high demand for storage.

Today, all storage access methods require CPU involvement in all I/O requests. With DAS, storage devices are very tightly coupled with the host computer's operating system and are typically managed using a parallel bus-based architecture such as SCSI. Storage sharing is limited because of its affinity (direct association) to the server. Additionally, these systems can pose maintenance intensive burdens as they require constant tuning in order to optimize CPU cycles between disk access and application processing.

Storage Area Networks

A storage area network (SAN) describes a dedicated, high-performance network infrastructure deployed between servers and storage resources (Figure 5). The storage area infrastructure is a separate, dedicated network entity optimized for the efficient movement of large amount of raw block data. In effect, SAN is an extended link between server and storage and enables the extension of the SCSI protocol over longer distances.

SANs are typically built using the SCSI and Fibre Channel (SCSI-FCP) protocols. Fibre Channel is well suited to this application because it can transfer large blocks of data (such as is possible with SCSI) while at the same time being able to transfer these blocks over longer distances (unlike SCSI). Fibre channel topologies, either loop or fabric, are built using specially designed networking gear that closely resemble the hubs, switches, and gateways used to build typical packet-based LANs and WANs.

The SAN market has historically addressed high-end, enterprise-class storage applications where performance, redundancy, and availability are paramount.

Figure 5 SAN storage model. Clients on the left communicate with the file servers on the right. Each server runs its own application and supplies its own associated storage. File server CPU must handle application traffic plus storage I/O requests, limiting their performance.

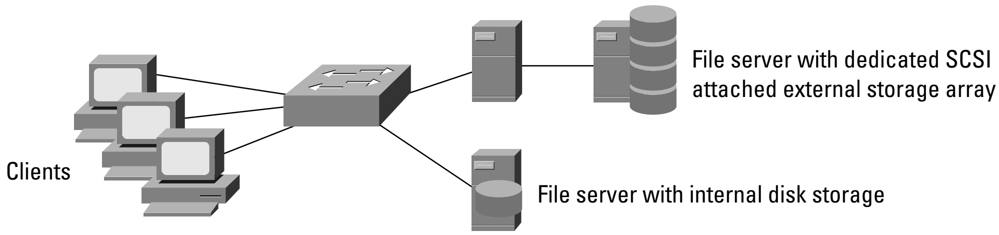

Network Attached Storage

Once LANs became a technological reality, it became possible to connect multiple file servers with a common infrastructure for the purpose of file sharing. LANs accelerated the demise of centralized computing and ushered in the development of distributed client/server computing. The concept of distributed client/server technology involved using an arrangement of inexpensive microcomputers and storage devices (disk, tape, and so forth) to replace the expensive centralized (mini, mainframe) computers.

As computers proliferated, there were many incompatibilities, which complicated data access. The advent and wide spread deployment of local area networking supported workgroup clusters with priorities for file sharing, interoperability and cost savings. Network-attached storage (NAS) consists of a specialized file server and storage (Figure 6). NAS servers run optimized file systems and install as preconfigured storage "appliances." Because NAS systems are connected to the LAN, clients are able to transfer data to and from the storage devices associated with the NAS system.

In addition, NAS can directly process file system protocols such as NFS and CIFS. Client machines "mount" volumes on these disk resources allowing virtual access to productivity applications and file data. While these devices are capable of hosting distributed applications, these applications typically are placed on application-specific file server platforms that have no storage I/O responsibilities.

Figure 6 NAS storage model. Clients on the left communicate with the application file server and NAS storage device.

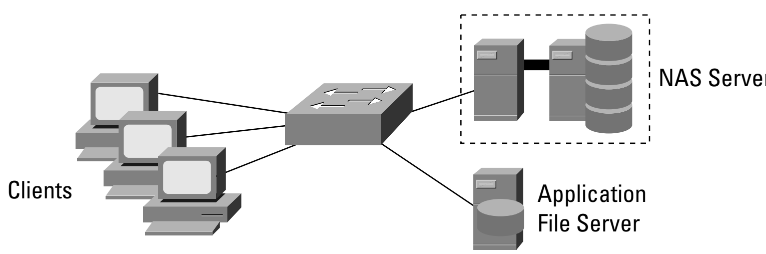

NAS systems generally support multiple operating systems and are directly accessible by heterogeneous clients via the LAN. Media can usually be replaced or added in these systems without impact to the network or application servers. NAS systems are highly suitable for most applications except high bandwidth video and larger, multi-user database applications, which will can bottleneck the LAN and slow the overall NAS appliance performance with TCP/IP overhead.

SAN Versus NAS: What Are the Differences and Similarities?

When it comes to determining the difference between SANs and NAS, many classify the two as competing technologies when they are in fact complimentary. SAN and NAS were independent events driven by different customer requirements. SAN is data centric whereas NAS is network centric. The following table better describes the functions and characteristics of NAS and SANs.

Table 1 NAS Versus Native Fibre Channel-Based SANs

Network Attached Storage	Storage Area Networks
Network centric	Data centric
LAN or WAN	Dedicated storage network
Defines a product/appliance	Defines an architecture
High latency—overhead of TCP/IP	Low latency—thin protocol
Indirect connection (routed)	Direct connection (switched)
Shared files, data	Server owns a volume

Source: IDC, Dataquest, Morgan Stanley Dean Witter Equity Research

NAS and SANs are complementary in that:

- NAS products could be placed on dedicated SAN networks to provide optimal performance for file protocol (NFS, CIFS) traffic.
- SANs can be expanded to include IP and other non-storage associated networking protocols.

The evolution of storage networking involves the fusion of the best features from DAS, NAS, and SAN, incorporated to create solutions which will meet the aggressive storage needs of today's Internet-centric business models.

TCP/IP Ethernet and Storage Networking

iSCSI

ISCSI is a new IETF draft that incorporates the "openness" of Ethernet & IP based NAS file-level access with the performance of SAN based block-level access. With the proliferation and speed of today's IP and Ethernet networks, customers can now take advantage of the same foundation that built the Internet (and the corresponding convergence of voice, video and data) to support the storage networking needs too. Servers located on Ethernet (100BaseT or GbE) running TCP/IP can load the open iSCSI driver and access block-level storage data located on local or remote fibre-channel SANs. Now customers can build storage networks using TCP/IP Ethernet yet still have the bandwidth and unlimited storage expansion capability. The Internet SCSI (iSCSI) protocol is a networking protocol for SCSI-3 traffic over TCP/IP Ethernet. It provides block-level I/O access to storage over the existing TCP/IP network. Due to the layering of the various SCSI protocols, the operating systems, file systems, and end-user applications are unaware of the actual transport supporting the storage I/O, be it SCSI, Fibre Channel or iSCSI. The interface and the functionality are thereby preserved.

iSCSI supports multiple topologies meeting specific requirements of IT professionals. It can be deployed as a private, storage centric network using well-known Ethernet devices with IPSec and firewalls for enhanced information security. It can be implemented as converged storage network, either as a private enterprise solution or public Internet (through virtual private networks or VPNs). Theoretically, customers can build networks of any size to support the diverse data demands of today's customers. It extends over unlimited distances all the while supporting the traffic of multiple users and applications within a single network infrastructure. Flexibility, scalability, security, management and storage consolidation can be achieved by combining the appropriate storage networking technology to create a total solution.

FCIP

FCIP (Fibre Channel over IP) is a proposed standard that provides fibre channel tunneling over TCP/IP. It involves the encapsulation of the Fibre Channel protocol within the IP packets so that it can be carried over the IP Network. Customers with pre-existing Fibre Channel networks can leverage their existing SAN and extend its reach over the WAN or MAN. FCIP is a viable means of connecting together Fibre Channel SAN "islands" which formed as a result of localized efforts of storage centralization.

For businesses that are expanding their storage infrastructures, they are faced with business continuance issues. Applications associated with disaster recovery and high availability can use FCIP as a solution for protecting their data. As an example, FCIP could be utilized to connect two geographically dispersed Fibre Channel storage arrays for the purpose synchronous data storage. If the local storage array was to become unavailable, an application could utilize the FCIP link to access the data on the "hot backup" storage system at the remote site. Also, remote tape backups can be implemented to further protect customer's valuable information in the event of a disaster at the primary site.

SAN Applications

One of the major motivators in moving to storage networks is the variety of applications that become possible with the implementation of this technology. These applications are capable of valuable performance enhancements, simplified storage management, and increased scalability to the storage infrastructure. An example of some these applications are:

- Data sharing—With storage devices centralized, multiple file servers can cost effectively access and share data without overall system performance degradation.
- Storage sharing—Two or more (homogeneous or heterogeneous) servers sharing a single storage entity that has been physically partitioned into sections, each only accessible by a specific server. It is possible for homogeneous servers to share a partition under certain circumstances.
- Data vaulting/data backup—In most scenarios, data back up / vaulting to near-line or off-line storage devices is dependent on the use of the common LAN or WAN infrastructure and its associated traffic characteristics. With SAN', these operations take place independent of the primary network, increasing the operational performance and enabling the use of new data-movement applications such as LAN-free or serverless backups.
- Disaster recovery—Traditionally, disasters involving data loss was handled using recovery from tape. The use of SANs allows for various methods of automated data backup and "warm" availability of this data in the event of a loss of the primary data image. Efficiently maintaining a redundant data image requires a low latency, high availability network infrastructure, which today's storage networks are suitable for. Total connectivity for the loss of the primary site would result in an image backup site emerging as the active image site. Recovery of the primary site is met with it being updated by the current active image and a decision whether to fall back to the original primary image server.

Storage Networking Futures

The active word in all of these scenarios is "networking." The distinction between DAS, SAS, SAN, and NAS is blurring. All technologies rise to the challenge of meeting storage demands of the end-user but some offer clear distinctions and benefits. The traditional client-server model of computing will evolve to a global storage network with any-to-any connectivity, enhanced data availability and optimal distributed data leveraging.

As with other leading edge technologies, the storage network marketplace will benefit from enhancements to existing technologies as well as the emergence of several new ones.

However, enhanced physical connectivity alone does not lift the barriers that keep storage resource from being easily shared by multiple servers. Advances in the storage network will take form in storage management, storage pooling, high availability, backup and restore automation, application integration, heterogeneous interoperability and quality of storage network services.

CISCO SYSTEMS

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy-les- Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 17, 99 Walker Street
North Sydney
NSW 2059 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

4/01 LW2320