A Critical Discussion of E-SAN and FC-SAN

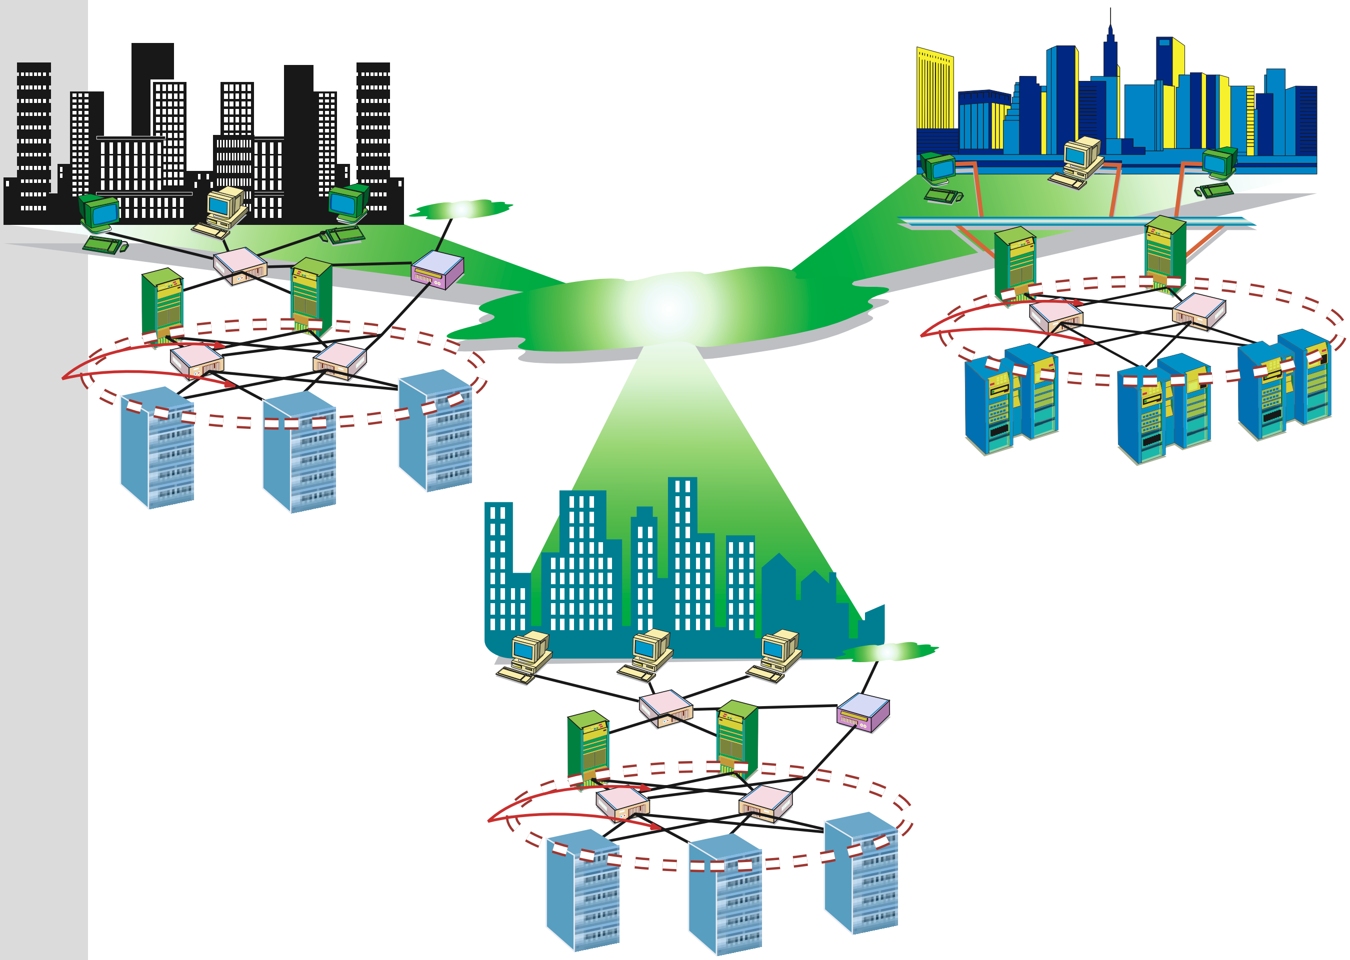

A business and technical perspective of Ethernet SANs (E-SANs) compared to Fibre Channel SANs (FC-SANs)

A
Critical
Discussion
of
E-SAN and FC-SAN

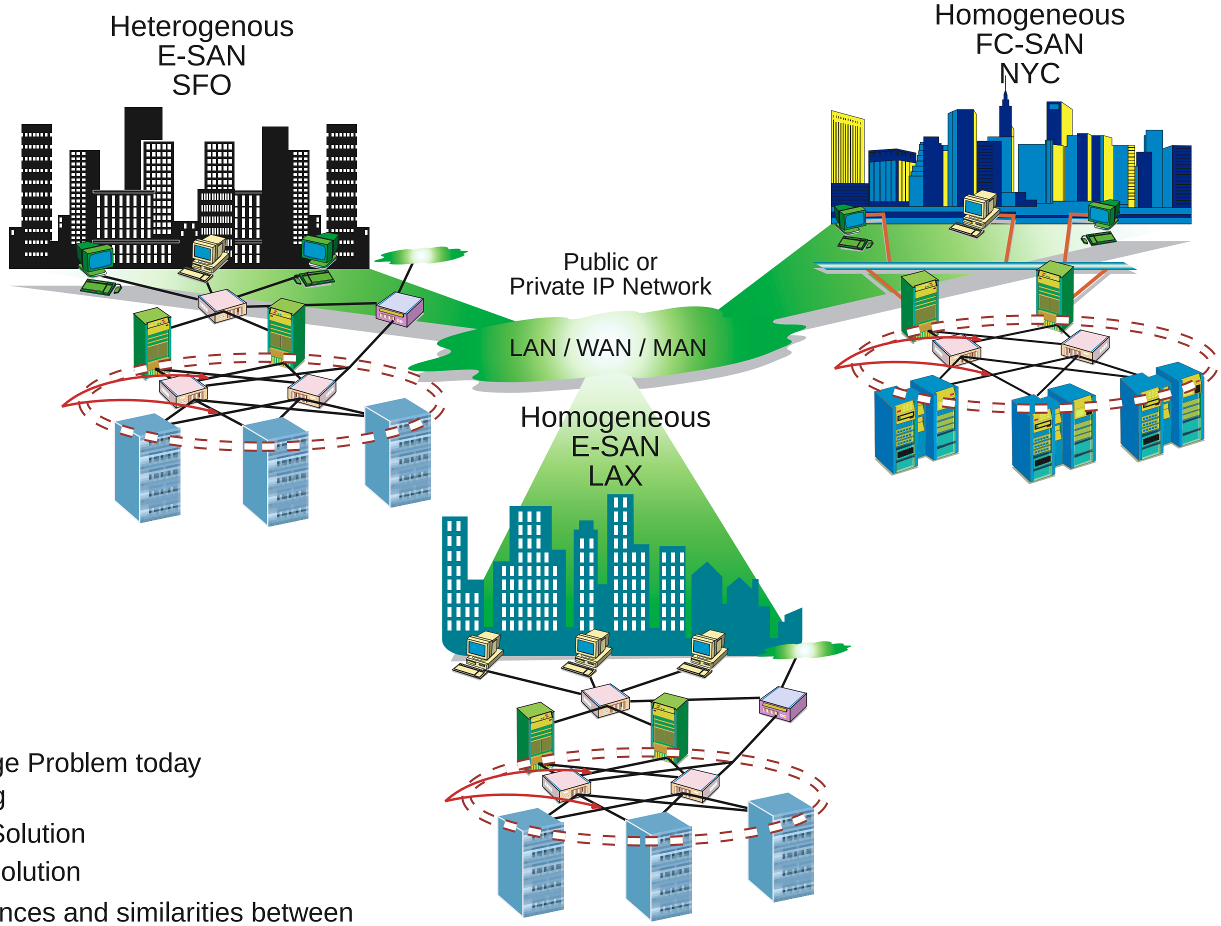

- The Storage Problem today
- Networking
- The SAN Solution
- The NAS solution
- The differences and similarities between NAS and SAN
- Fibre Channel (What is it?) and FC Storage Area Networks (FC-SANs)
- Ethernet (What is it?) and Ethernet Storage Area Networks (E-SAN)
- The Advantages of E-SANs over FC-SANs
- 3ware's Network Storage Unit (NSU)
- A summary discussion of trends and predictions into the future

Preface

Situation Overview

The topic of this report is really the story of the storage industry undertaking to develop an entirely new networking standard (Fibre Channel) with the goals of providing an optimized environment for centralized data management and separation of storage and message traffic on the network. The implementation of FC is not complete in the areas of security, network management, congestion control, and ease of deployment. In addition, the technology requires significant investment and the learning of new skills. In fact this goes against the trend to minimize operations complexity by focusing on "best of breed" technologies. From its very beginning, Fibre Channel has been influenced by the desire of storage vendors to optimize their available technology. The result has been to create an environment in the Fibre Channel world that is expensive to implement and does not offer the interoperability requested by the IT community.

FC is not complete in the areas of security, congestion control, and ease of deployment.

Because of the issues identified above, visionary companies like 3ware recognized that Fibre Channel could be limited to the data center and reopened the topic of how the benefits of IP networking could be used to make the SAN vision a reality. It became obvious with the advent of gigabit Ethernet that all of the advantages of the Fibre Channel SAN vision could be delivered using TCP/IP while side stepping many of the deterrents to FC-SAN that are now painfully obvious from real-world experiences in the marketplace. Instead of requiring customers to buy expensive FC adapters, optical fibre cabling, new FC switches and new network analyzer tools and training, E-SANs could use commodity Ethernet adapters, CAT5 or optical Fibre cabling and switches.

3ware reopened the topic of how IP networking could be used to make the SAN vision a reality.

To do this however required that the performance issues of TCP/IP stack processing be overcome. There are a variety of approaches today to solve these issues and 3ware believes that their innovative storage protocol SCP and internal packet switching architecture is a correct approach. The trend to outboard TCP/IP stack processing among all Ethernet NIC vendors, and the upcoming speed advantage that 10 Gigabit Ethernet will enjoy over next generation Fibre Channel makes a strong case for E-SAN and the 3ware NSU architecture.

Table of Contents

Table of Figures

Tables

The Storage Problem Today

1

The driving need to consolidate and centrally manage data

Storage is predicted by industry analysts to account for over 75% of all computer hardware expenditures over the next five years. It is not uncommon for data to be scattered over numerous and distributed UNIX and NT servers in a typical enterprise following the traditional client-server model. It is increasingly important for both large and small enterprises to have a consistent and effective strategy regarding storage architecture and data management. SAN is a solution that accomplishes this objective by providing centrally managed storage throughout the enterprise while at the same time providing high speed backups that address the number-one reported storage administration issue today - backup windows.

Not only do SANs have the advantages of reducing operating costs by centralizing data management functions, but they also improve performance and availability by consolidating environments where user data could be spread over 100 or more application servers. In highly decentralized environments, server node failures could lock users out of critical data that is not adequately protected. This problem is easier to solve with SANs which can be redundantly configured to prevent unplanned interruptions in data accessibility.

SANs have the advantages of reducing operating costs by centralizing data management functions.

The importance of understanding networks and network terminology

Those readers who already understand the basic terms and concepts of networking can skip Chapters 2 and 3. An analogy we will use in this chapter to explain networking terminology is to compare sending a message over a network to sending a letter through the U.S. post office from one point to another throughout the world. Local Area Networks (LANs), Wide Area Networks (WAN), networking cable and equipment, the Internet, many types of servers, and the different types of computer storage equipment comprise the basis for the analogy.

A standard way of viewing the Network

For consistency throughout this paper we will use a standard top to bottom scheme for the different layers of protocol through which messages pass as they go from their source to their destination. In Figure 1, the “Source” is shown at the top left and the “Destination” is shown at the top right. Levels of protocol from the application layer to the physical layer are shown top to bottom as illustrated in Figure 2.

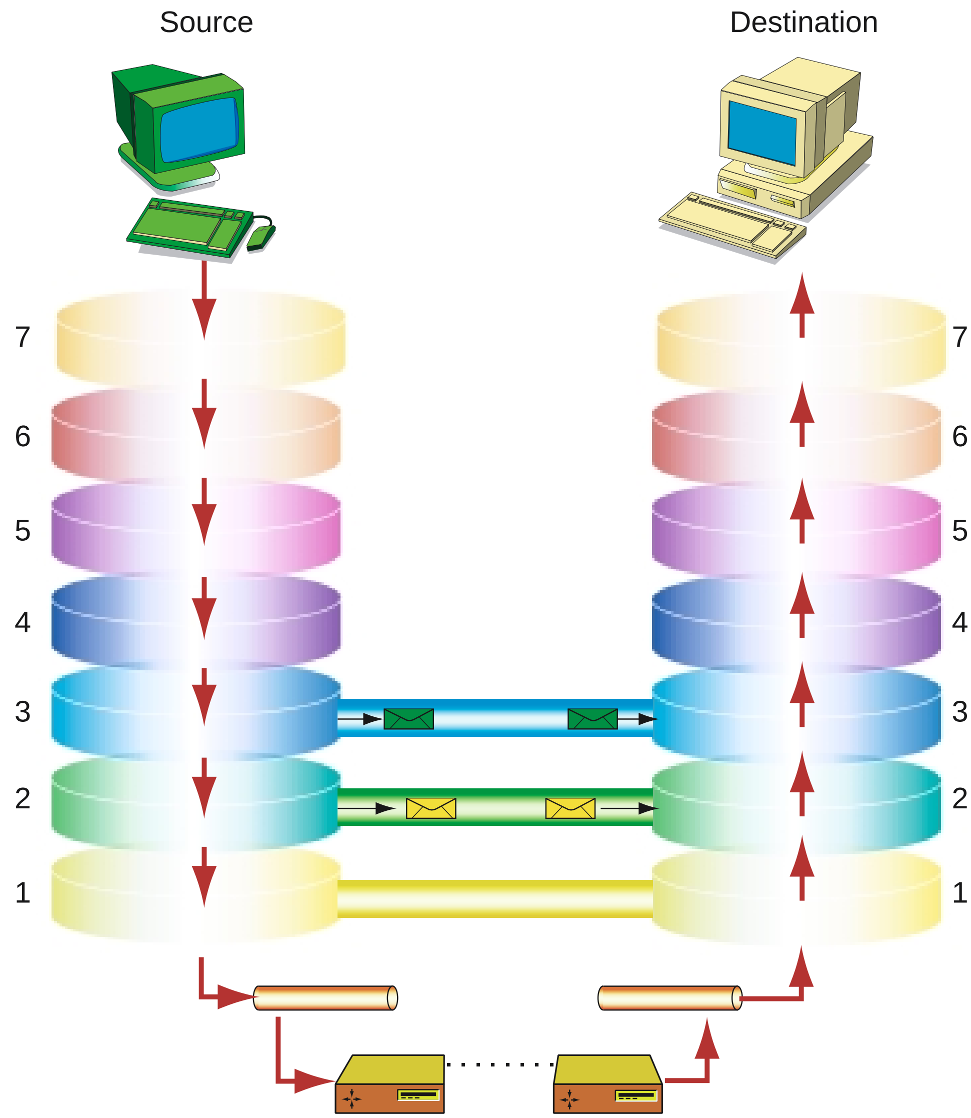

Figure 1 - A standard way of viewing the networking model

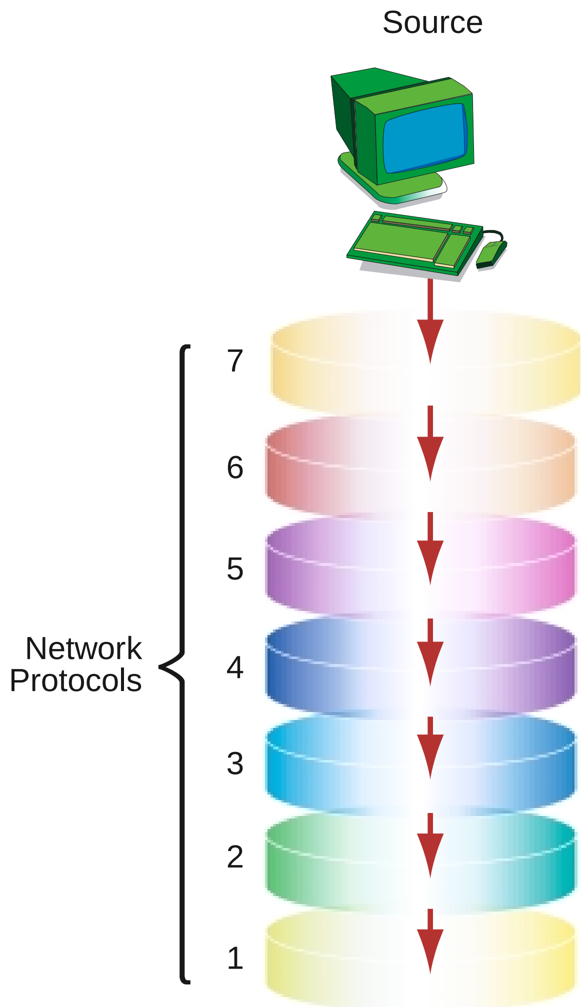

Figure 2 - A standard way of viewing networking protocols

Other Information Sources

This paper augments the concepts discussed in the 3ware white paper ***Gigabit Ethernet SANs with switched RAID arrays*** which is available at www.3ware.com. This report also contains a complete Glossary of Terms and many Figures and Tables that attempt to clearly explain the technical topics involved in networking.

Networking

2

The OSI reference model defined

Networking protocols are typically described in reference to the OSI model, which defines the different layers through which a message passes from one computer to another. OSI stands for Open System Interconnection, an ISO standard for worldwide communications that defines a networking framework for implementing protocols in seven layers. Control is passed from one layer to the next, starting at the application layer in one station, proceeding to the bottom layer, over the channel to the next station and back up the hierarchy. Although the OSI model was never fully implemented as a standard, it defines the functionality through which all messages pass in the communications process and is the teaching model for all networking today. Most of the functionality in the OSI model exists in all communications systems, although two or three OSI layers may sometimes be incorporated into one. Table 1 and Figure 3 describe the seven layers of the OSI model:

The OSI model defines the functionality through which messages pass in the communications process.

Level or Layer	Name	Function
7	Application Layer	Program-to-program communication.
6	Presentation Layer	Manages data representation conversions. For example, the Presentation Layer would be responsible for converting from EBCDIC (IBM mainframes) to ASCII (UNIX and Windows NT).
5	Session Layer	Responsible for establishing and maintaining communications channels. In practice, this layer is often combined with the Transport Layer.
4	Transport Layer	Responsible for end-to-end integrity of data transmission.
3	Network Layer	Routes data from one node to another, and fragments and reassembles packets to match the constraints of the specific networks
2	Data Link Layer	Responsible for physically passing data from one node to another.
1	Physical Layer	Manages putting data onto the network media and taking the data off.

Table 1 - The OSI Reference Model structures communication software in layers

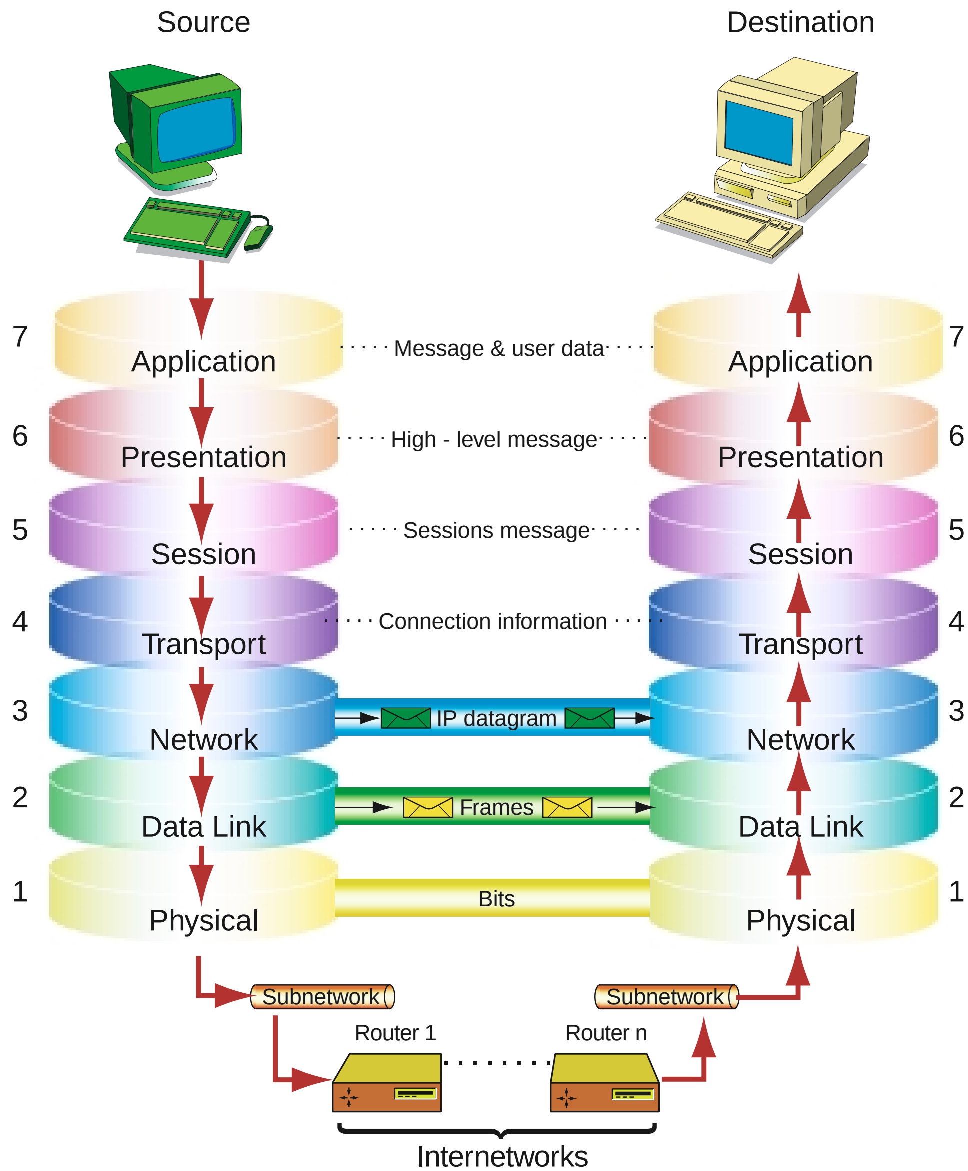

Figure 3 – OSI Reference Model structures communication software in layers

Today' s network protocols are many and varied. Understanding where a particular protocol resides is shown in Table 2 and Figure 4. Efficient communications in a network depend on choosing the correct protocol for each level of the networking process so that the services that are provided are appropriate for the type of data in the message. A simple example of this would be the difference between UDP and TCP. UDP is used in time sensitive audio and video broadcasting and drops frames in the event of a problem. Missing some pixels in a video transmission of course is less important than missing the continuity of a live real time broadcast. TCP however continues to resend the frames until they are successfully delivered.

Level or Layer	Name	Examples of Network applications and protocols (See the glossary for definitions)
7	Application Layer	In the Internet world common applications include Telnet, FTP, HTTP, NTTP, SHTTP, IRC, NFS, RPC, CIFS, SMTP and DNS
6	Presentation Layer	EBCDIC, ASCII, XDR
5	Session Layer	SEP, sockets
4	Transport Layer	TCP, FCP, SPX, STP, UDP, iSCSI[1], IPFC
3	Network Layer	IP, IPX, CLNP, FC
2	Data Link Layer	PPP, SLIP, FC, LAP protocols i.e., MLP,LAPD, LAPM, LAPF
1	Physical Layer	10Base T, 100Base T, 1000Base T, 10,000Base T, ATM, ESCON, FDDI, FICON, HIPPI, RS 232, RS 449, FC

Table 2 - Today's network protocols as they relate to the OSI model

Figure 4 shows the function of each protocol layer of the OSI model and shows the way that subnetworks are connected by routers into internetworks that pass the message from source to receiver. Within a subnetwork, switches filter and forward packets between Local Area Network (LAN) segments. Switches operate at the data link layer (Layer2) of the OSI Reference Model and therefore support any packet protocol.

Switches filter and forward packets between Local Area Network (LAN) segments.

Special types of switches called L3 and L4 switches operate at the third (Network Layer) and fourth layer (Transport Layer) of the OSI model and are responsible for the integrity of data transmissions between LAN segments. LANs that use switches to join segments are called switched LANs or in the case of Ethernet networks, switched Ethernet LANs.

[1] iSCSI is not yet a protocol, since it has no agreed spec. There is not yet consensus on whether it is on top of TCP, IP or even STP. It will probably be a layer 4 protocol as shown in Table 2.

Figure 4 – Network protocols as they relate to the OSI model

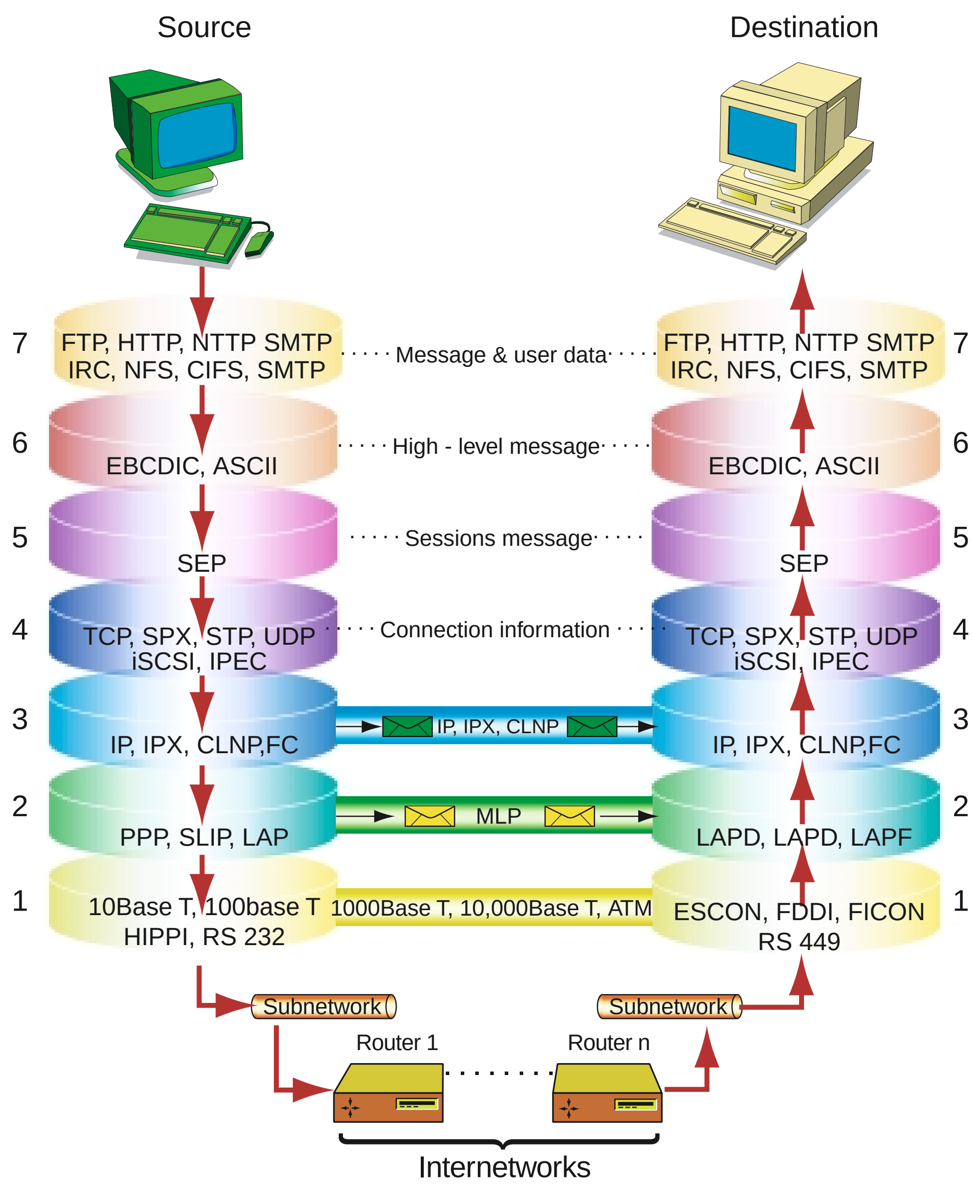

Comparison of the OSI and Internet Networking Models

The Internet Protocol (IP) model collapses the top two layers into one but the functionality is the same as for the OSI model. The Internet is going to be increasingly popular for storing and retrieving data - especially block oriented data. This is the interest of the IETF group on IP storage that is discussed further in Chapters 6 and 7. A comparison of the OSI and Internet models is shown in Figure 5.

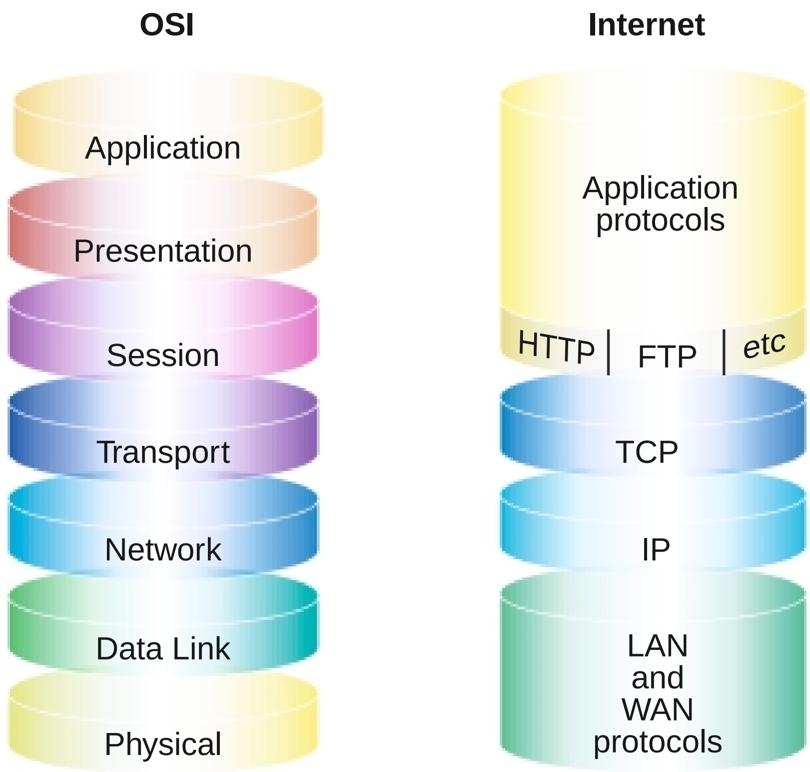

Figure 5 - A comparison of the OSI and Internet protocol stacks

Backbones

Backbones are connected by routers.

As shown in Figure 6, the main function of backbones is the sending of frames from a source system to a destination system. This is done over backbones. Backbones are connected by routers that examine IP envelope addresses and send them on their way to the correct destination. Some backbones are paid for by government agencies such as the National Science Foundation (NSF) or the National Aeronautics and Space Administration (NASA). Backbones connect Network Access Points (NAPs) throughout the world and many are based on ATM OC-12 at 622 Megabits per second (Mbps). A future higher speed backbone known as vBNS (very high-speed Backbone Network Service) is in limited deployment at this time and runs at 2.5 Gigabits per second (Gbps). Local networks and Web Sites are connected to the Internet in a variety of ways. A residential telephone line can transmit data at 56kbps (kilobits per second), and higher bandwidth leased telephone lines can carry data at 1.544Mbps for T1 links and 44.746Mbps for T3 links.

Regional networks

Regional networks maintain and provide Internet access within a certain geographic area with a worldwide presence.

Routers

In networking, frames are like postal envelopes with addresses called headers. Routers send them from one subnetwork to another to their final address. Along the way routers look at the header [address] and forward the frame [letter] onward to its destination much like automated postal scanning equipment scans a zip code and forwards it to the correct post office for delivery.

Routers connect backbones and do most of the work directing traffic over the Internet.

Routers connect backbones and do most of the work directing traffic over the Internet. When TCP creates a packet it adds a header to the data which defines how the packets are to be reassembled at the other end. It also adds a checksum to determine if any errors occur during transmission. The packets are then placed into one or more IP envelopes, which add another type of header that contains information about the address to which the packet is to be sent. Using standard protocols such as Routing Information Protocol (RIP) for interior routers and Exterior Gateway Protocol (EGP) for exterior routers, the router examines the header for each IP envelope to determine the exact IP address to which the envelope is to be sent. This process is illustrated in Figure 6.

Figure 6 - Routers send frames with encapsulated packets across internetworks

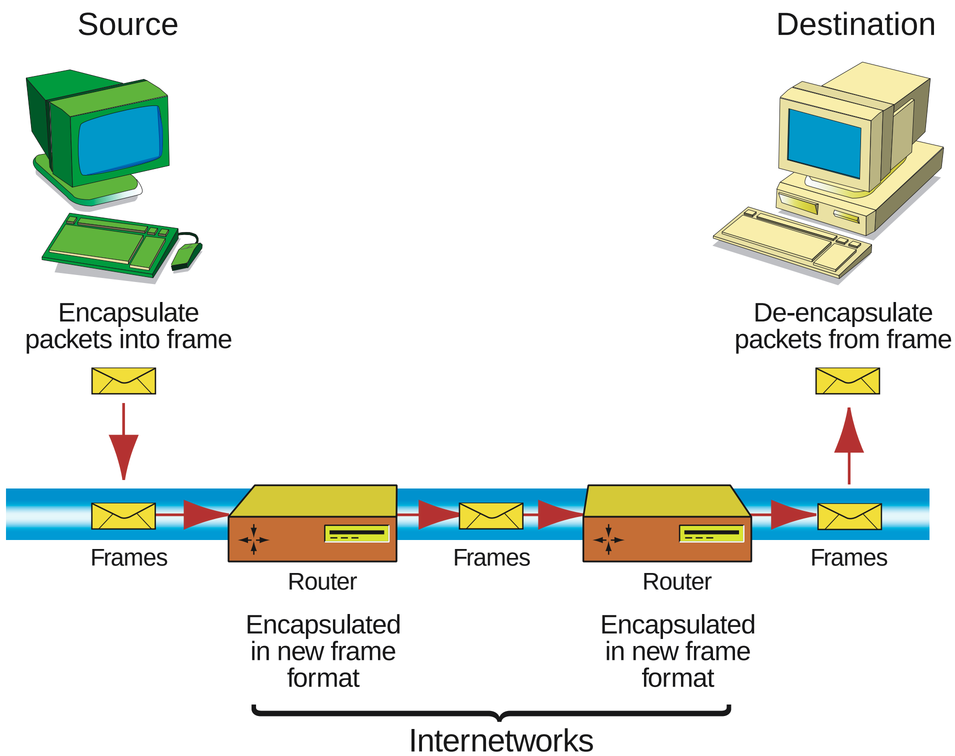

Bridges, Hubs and Switches

A bridge is a device that connects two local-area networks (LANs).

A bridge is a device that connects two local-area networks (LANs), or two segments of the same LAN. The two LANs being connected can be alike or dissimilar. For example, a bridge can connect an Ethernet with a Token-Ring network. Unlike routers, bridges are protocol independent. They simply forward packets without analyzing and re-routing messages. Consequently, they are faster than routers, but also less versatile.

Hubs and switches provide a common connection point for devices in a network. Hubs are commonly used to connect segments of a LAN. A hub contains multiple ports. When a packet arrives at one port, it is copied to the other ports so that all segments of the LAN can see all packets. Switches also connect multiple LAN segments, but they connect only the input port to its desired output port instead of broadcasting incoming packets to all outputs. A hub has less performance than a switch, because it does not support multiple simultaneous connections between pairs of ports. Hubs are not supported with Gigabit Ethernet.

TCP/IP

Transmission Control Protocol (TCP) breaks down and reassembles packets that are sent over the Internet and Internet Protocol (IP) makes sure that packets are sent to the right destination. TCP and IP therefore are the two most important communications protocols on the Internet. They are frequently referred to as TCP/IP. Packets can travel over the Internet on separate routes and TCP will keep re-sending packets if they are lost or if other problems occur. “Packet loss” is one of the biggest reasons for delays in message transmission. The software that is required for a computer to understand and execute TCP/IP protocol is known as a TCP/IP stack. This software acts as an intermediary between the network and a computer.

TCP and IP therefore are the two most important communications protocols on the Internet.

HTTP

HTTP is an abbreviation for HyperText Transfer Protocol, the underlying protocol used by the World Wide Web. HTTP defines how messages are formatted and transmitted, and what action Web servers and browsers should take in response to various commands. For example, when you enter a URL in your Web browser, this actually sends an HTTP command to the Web server directing it to fetch and transmit the requested Web page.

HTTP is the underlying protocol used by the World Wide Web.

UDP/IP

User Datagram Protocol (UDP) is used instead of TCP for certain types of data such as audio data. UDP is a simpler protocol than TCP and does not guarantee data delivery. It is most suitable for applications, such as audio, where real-time delivery is more important than assuring that all data is delivered.

FTP

The Internet’s File Transfer Protocol or FTP is the most common protocol used for downloading files to your computer or uploading files to another computer. The files can vary from programs that can run on a computer to graphics that can be viewed or sounds and music that can be heard. Special software is required to connect to an FTP server, which is available from many places on the Internet.

SLIP and PPP

Serial Line Internet Protocol (SLIP) or Point-to-point (PPP) protocols are methods of connecting a computer to the Internet via a serial port and a modem when a network card is not used.

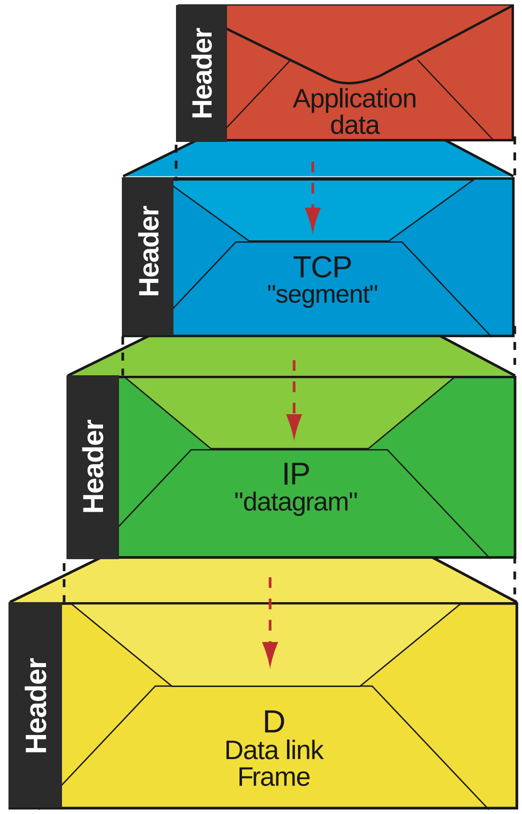

Figure 7 - Frames contain "nested" headers and trailers for each level of the OSI model

Files types of the application layer

Files have different formatting parameters.

There are a variety of file types with specific formatting and layout parameters that can be packaged in IP envelopes and sent over a network to other computers. The other computer must have a program called a "reader" that allows the particular file format to be displayed on the screen or activate video cards, sound cards or other devices.

Table 3 lists popular file types that can be sent over the Internet.

File type	File extension	Definition
BMP	.bmp	Standard bitmap Windows graphics format
Excel	.xls	Microsoft Excel file format
EXE	.exe	Executes programs
GIFF	.gif	Standard graphics file for the World Wide Web
JPEG	.jpg	Compression standard for color images
MIME	.mim	A standard for non-ASCII messages that allows them to be sent over the Internet
PCX	.pcx	Graphics file format for PCs, scanners, fax
PDF	.pdf	Adobe's Portable Document Format
PNG	.png	A license free graphics format similar to GIFF
PKZIP	.zip	Popular shareware compression technique
Powerpoint	.ppt	Microsoft PowerPoint file format
TIFF	.tif	A bitmap graphics format for Apple and PCs
WAV	.wav	Microsoft standard for sound files
Word	.doc	Microsoft Word file format

The SAN Solution

3

What is SAN

SANs are subdivided into FC-SANs and E-SANs.

There is much confusion today about Direct Attached Storage (DAS), Network Attached Storage (NAS) and Storage Area Networks (SAN) which comprise the three major types of storage architectures today. When to use each in enterprise storage architecture is often misunderstood among storage architects. SANs are further subdivided into Fibre Channel SANs (FC-SANs) and Ethernet SANs (E-SANs) depending on the type of network used.

The major differences between DAS, NAS and SAN depend on

1) Whether or not a network is involved in the I/O path

2) The placement of the network relative to the file system

DAS can be explained as a three-element model and NAS, FC-SAN and E-SAN as four-element models. This is shown in Figure 8 below.

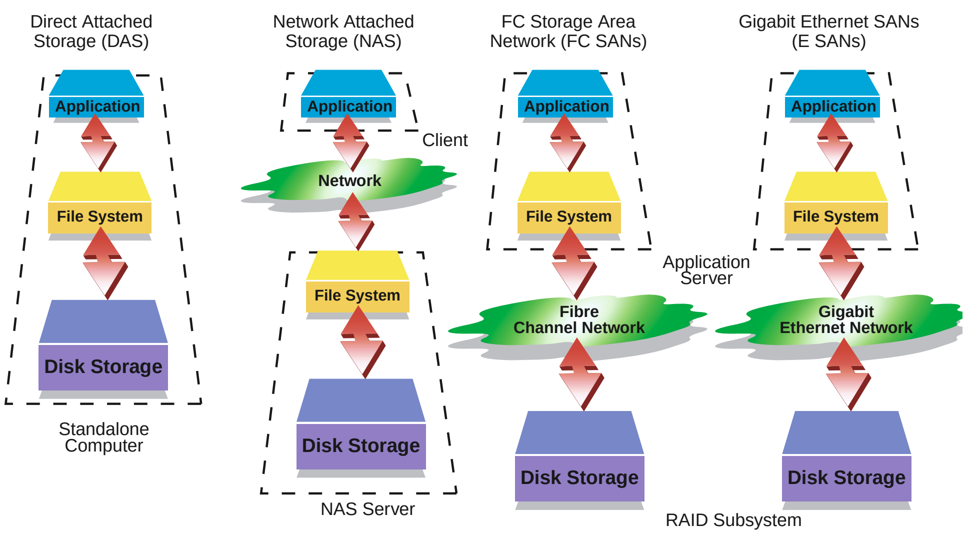

Figure 8 – There is much confusion today about DAS, NAS and SAN and when to use each in an enterprise storage architecture

The benefits of SAN

To understand the benefits of SAN it is important to understand the differences between DAS, and SAN as shown in Table 4. NAS and SAN comparisons are discussed in chapter 4.

Table 4 - The differences between DAS and SAN

Storage characteristic	DAS	SAN
Sharing storage across heterogeneous OS	No	Yes
Data security	Server dependent	Server, switch and storage dependent
Installation	Varies	FC-SAN - difficult E-SAN - easy
Effective centralized management	Varies	FC-SAN - difficult E-SAN - easy
Suitable for deploying many distributed clients	No	Yes
Server independent backup	No	Yes
Disaster tolerance	Yes	Yes
Server free backup	No	Yes
Highly scaleable	No	Yes

The differences between FC-SANs and Ethernet SANs

FC-SAN and E-SAN uses different transport and network protocols.

The major difference between FC-SAN and E-SAN is that FC-SAN uses different transport and network protocol (FCP/FC) for the transmission of data compared to E-SAN, which uses TCP/IP. This is shown in Figure 9.

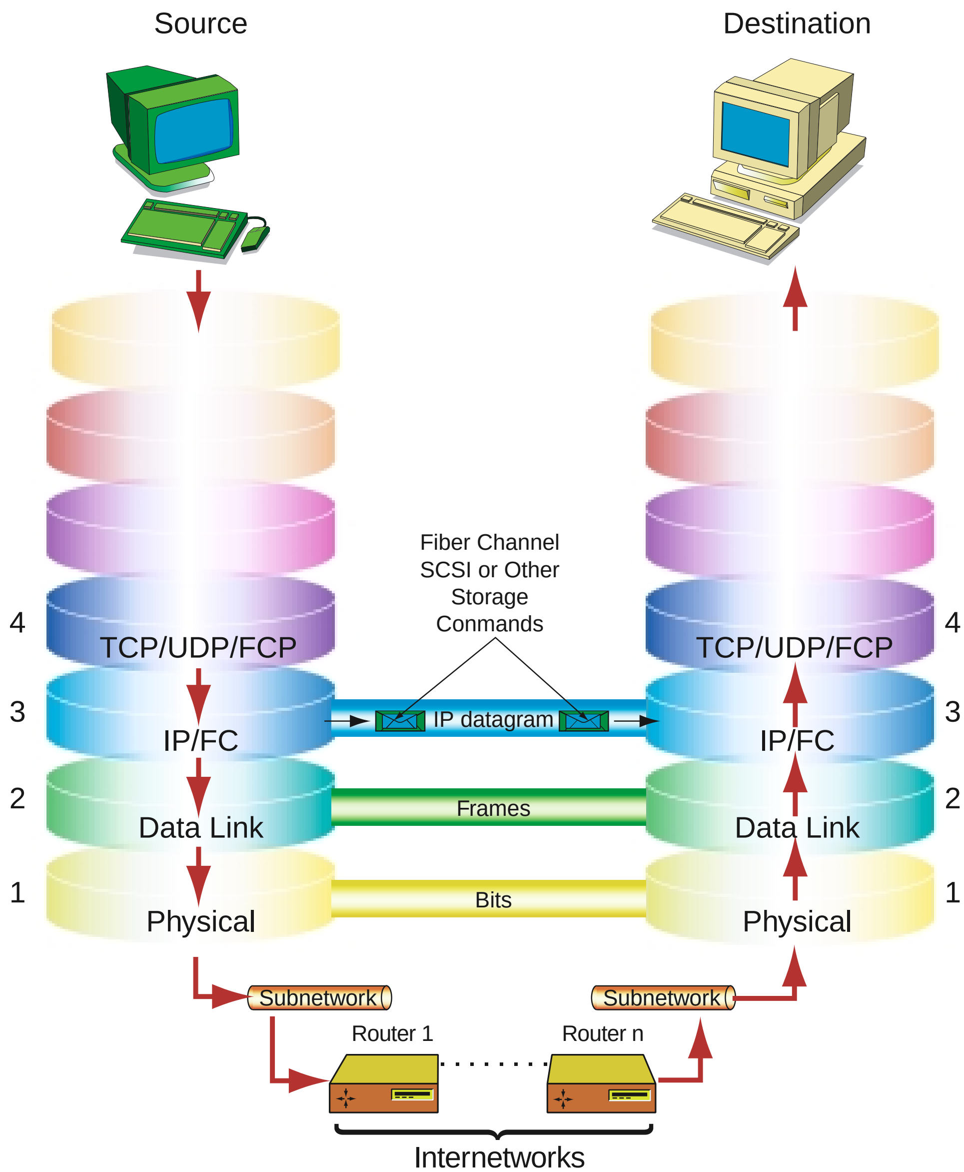

Figure 9 - Fibre Channel SANs use FCP/FC while E-SAN uses TCP/IP for transport and network protocols

FC-SAN and E-SAN are two very different technologies.

There are many differences between FC-SANs and E-SANs, which are discussed further in Chapters 7 and 9 of this report. The major differences are shown in Table 5 with the relative advantages of each technology highlighted in blue. It is very clear that FC-SAN and E-SAN are two very different technologies. There are major issues such as cost, “proprietary lock-in”, lack of standards and lack of training deterrents to Fibre Channel SANs that are not deterrents to E-SAN. These problems have slowed the adoption rate of FC-SAN, keeping it at about 1 in 10 enterprises today. Perhaps the most important advantage of FC-SAN is that it provides the data center with a solution today. Because of its size and importance. FC-SAN can be better justified in this environment than in the distributed departmental environment. The 3ware NSU is the first available storage product that works in an E-SAN, to extend the benefits of SAN beyond the data center.

Looking into the future, network speed for the next generation technologies is another major advantage for Ethernet. Fibre Channel has been specified at a

2 Gigabits/sec performance compared to Ethernet which has been specified to at 10 Gigabits/sec. Add to this the 10-kilometer distance limitation of Fibre Channel compared to the virtually unlimited distances of E-SAN and the case becomes more compelling for E-SAN. Considering further the broadly accepted standards for security and congestion control of E-SAN and it is very likely that functions that do not implement well in SANs today, will wait for the emergence of E-SAN and the many advantages it brings in interoperability, security, manageability and cost effectiveness.

Table 5 - The major differences between FC-SANs and E-SANs

Type of SAN	FC-SAN	E-SAN
Current generation bandwidth	1 Gigabit/sec	1 Gigabit/sec
Next generation bandwidth	2 Gigabit/sec	10 Gigabit/sec
Current deployment	8-10% of enterprises	Emerging technology Deployments expected in 2001
Distance supported	10 Kilometers	Nearly unlimited distance - but practical limits of 10 km.
Networking Protocols	Fibre Channel has a protocol for each layer of the OSI model with (FCP/FC) for the transport and network layers.	TCP/IP
Storage Protocol	Fibre Channel	SCSI or other protocols such as 3ware's Storage Control Protocol (SCP)
Storage Standards	FC protocol	3ware SCP. An iSCSI model is emerging from IETF.
Integration effort	Integrators are inexperienced and there is much difficulty with most FC-SAN deployments. Vendors have extensive interoperability labs that must test all components.	Because of the experience base and maturity of network standards that exists, E-SAN is easy to configure and integrate.
Data security	There are Fibre Channel security risks due to lack of standards	Ethernet provides well understood security protection through well known Network facilities

The current confusion over which approach to SAN is best.

E-SAN has major advantages over F-SAN.

There is currently a significant debate in the storage and networking communities over whether FC-SAN or E-SAN is best for a particular application. Although E-SAN is new and not deployed yet, 3ware believes that E-SAN provides major advantages of ease in installation, affordability, interoperability and long distance connectivity, standards, use of existing skills and affordability, and will quickly gain significant support. Especially important is the fact that the 10Gig Ethernet will be deployed before Fibre Channel reaches that performance level. Fibre Channel adoption of 10 Gbit links will be delayed because many vendors will push the intermediate step to 2 Gigabits per second links. Chapters 6, 7, 8 and 10 discuss this topic in detail.

The NAS Solution

4

What is NAS?

The NAS model was made possible by the fact that NFS (for UNIX) or CIFS (for Windows) allows a filesystem to be located (mounted) remotely and accessed over a network instead of residing on the application server. In the case of the NAS model, the application software makes a request for I/O to the remote filesystem mounted on a NAS server. The filesystem on the NAS server in turn determines the location of the data that has been requested by the application on the client. If the data is not in the NAS filesystem cache, the NAS filesystem then makes a request to the disk controller that retrieves the data from its disks array and returns the data to the NAS filesystem, which sends the data for the file to the client across the network. Compared to DAS, NAS servers offload some of the functions of organizing and accessing all directories and data on disk and managing cache. In theory, this frees the server's CPU to do additional work thereby reducing potential CPU bottlenecks. In practice, the actual CPU consumption depends on the efficiency of the NFS and TCP implementations compared to the local filesystem implementation.

NAS servers offload some of the functions of organizing and accessing all directories and data on disk.

Pros and Cons of Network Attached Storage

Network Attached Storage is not appropriate for all applications since not all applications use filesystems. These applications include some Databases where the storage device is mounted as raw storage thereby bypassing the filesystem. NAS devices present the data to the applications as files. DAS and SAN devices present the data to the applications as data blocks. An overview of the pros and cons of NAS is shown in Table 6 below.

Network Attached Storage (NAS)	
Pros	**Cons**
- Easy to Use - True File Sharing - Heterogenous Interoperability - Networking Standards • NFS • CIFS - Uses Existing Gb/Ethernet Infrastructure - Server Free and LAN free backup	- File Transfers only - Not Suitable for applications requiring blocks • Data Base • OLTP - Requires Distributed - File Management System - The filesystem makes price more like a server than storage

Table 6 – The Pros and Cons of NAS

The Differences and Similarities Between NAS and SAN

5

A comparison of SAN and NAS

There are many common characteristics of NAS and SAN because of the both architectures connect storage over the network. In fact, the two topologies are more complementary than competitive. Because of the outboard file system in NAS products, they are preferred for many applications that require UNIX and NT single image secure file sharing from a common pool of storage. This is because the techniques to share data over SANs are more complex and require emulation by NT clients for UNIX remote file sharing and vice versa. In contrast, NAS products contain a file system which handles all permissions for either UNIX (NFS) or Windows (CIFS) locally.

Common benefits of SAN and NAS are shown in Table 7.

Common Benefits of NAS and SAN Storage Architectures	NAS	SAN
Remote connectivity of storage	Yes	Yes
Reduced storage management costs through consolidation	Yes	Yes
Ease of scaling (adding) storage capacity and processor performance	Sometimes	Yes
Separation of storage and server purchases	Sometimes	Yes
Centralization of storage management and support	Yes	Yes
Single image data availability to users with different operating systems	Yes	Difficult
Removal of backup data flow from the LAN	Sometimes	Yes
Server-free backup	Yes	Yes
Increased storage resource utilization	Yes	Yes

Table 7 - The common characteristics of NAS and SAN

Fibre Channel (What is it?) and FC Storage Area Networks (FC-SANs)

6

What is Fibre Channel

Fibre Channel SANs are based on **Fibre Channel (FC)** which is an ANSI standard designed to provide high-speed data transfers between workstations, servers, desktop computers and peripherals. Fibre Channel makes use of a circuit/packet switched topology capable of providing multiple simultaneous point-to-point connections between devices. The technology has gained interest as a channel for the attachment of storage devices, but has limited popularity as a high-speed network interconnect. Fibre Channel can be deployed in point-to-point, arbitrated loop (**FC-AL**), or switched topologies. Fibre Channel **nodes** log in with each other and the switch to exchange operating information on attributes and characteristics. This information includes **port names** and **port IDs** and is used to establish interoperability parameters. Many of the switch vendors are still working out the bugs in Fibre Channel switches and as recently as June 2000 a senior Brocade executive (Brocade is a major FC switch vendor) admitted to a Gartner Group audience that the Fibre Channel fabric "was not there yet"

Fibre Channel has limited popularity as a high-speed network interconnect.

EMC's FC-SAN is called Enterprise Storage Networks. EMC is currently a member of the Storage Networking Industry Association (SNIA) but has also created and supported the Fibre Alliance, which is developing a rival set of standards.

Compaq's Enterprise Network Storage Architecture (ENSA) is another major proprietary FC-SAN that was announced in 1999 and is still largely a vision.

The SNIA standards efforts for FC-SANs

The Storage Networking Industry Association (SNIA) is a not-for-profit organization, made up of 150 companies and individuals spanning virtually the entire storage industry. SNIA members share a common goal, to set the pace of the industry by ensuring that storage networks become efficient, complete, and trusted solutions across the IT community. To this end the SNIA is committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. The SNIA is working on standards for Fibre Channel SANs and for Network Data Management Protocol NDMP.

Fibre Channel Storage Area Networks (FC-SANs) defined

Historically in storage environments, physical interfaces to storage consisted of parallel SCSI channels supporting a small number of SCSI devices. With Fibre Channel, the technology provides a means to implement robust storage area networks that may consist of hundreds of devices up to an architectural limit of 16 million. Fibre Channel Storage Area Networks (FC-SANs) yield a capability that supports high bandwidth storage traffic on the order of 100 MB/s, and enhancements to the Fibre Channel standard will support 200 MB/sec by 2001-2002.

Depending on the implementation, several different topologies components can be used to build a Fibre Channel storage area network. The Fibre Channel SAN consists of components such as storage subsystems, switches, and server systems that are attached to a Fibre Channel network using Fibre Channel host bus adapters. Fibre Channel networks in turn may be composed of many different types of interconnect entities. Examples of interconnect entities are switches, hubs, and bridges. As different types of interconnect entities allow Fibre Channel networks to be built of varying scale. In smaller SAN environments, Fibre Channel arbitrated loop topologies employ hub and bridge products. Larger networks would incorporate switch technologies.

Each of the components that compose a Fibre Channel SAN must provide an individual management capability.

As SANs increase in size and complexity, more Fibre Channel switches will be introduced. Each of the components that compose a Fibre Channel SAN must provide an individual management capability, and participate in a complex management environment. Due to the varying scale of SAN implementations, it is useful to view a SAN from both a physical and logical standpoint. Similarly, the logical view allows the relationships and associations between SAN entities to be identified and understood.

Because of the lack of standards for interoperability between Fibre Channel switches from different vendors, implementations of FC-SANs today tend to force and "lock-in" an enterprise into using a homogeneous FC-SAN with all switch and storage elements of the SAN from a single vendor. This is shown in Figure 10 below.

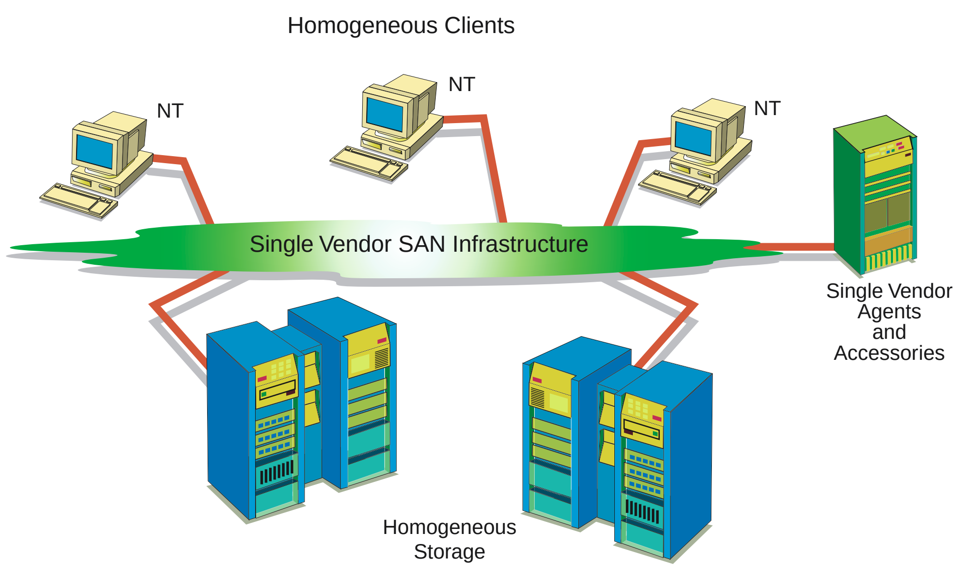

Figure 10 - A homogeneous Fibre Channel Storage Area Network or FC-SAN

The SAN vision of heterogeneous clients, switches and storage systems therefore is far from reality. For this reason, it is possible that E-SAN [although later to market] may well be deployed in heterogeneous configurations before FC-SAN. The heterogeneous vision for FC-SAN is shown in Figure 11.

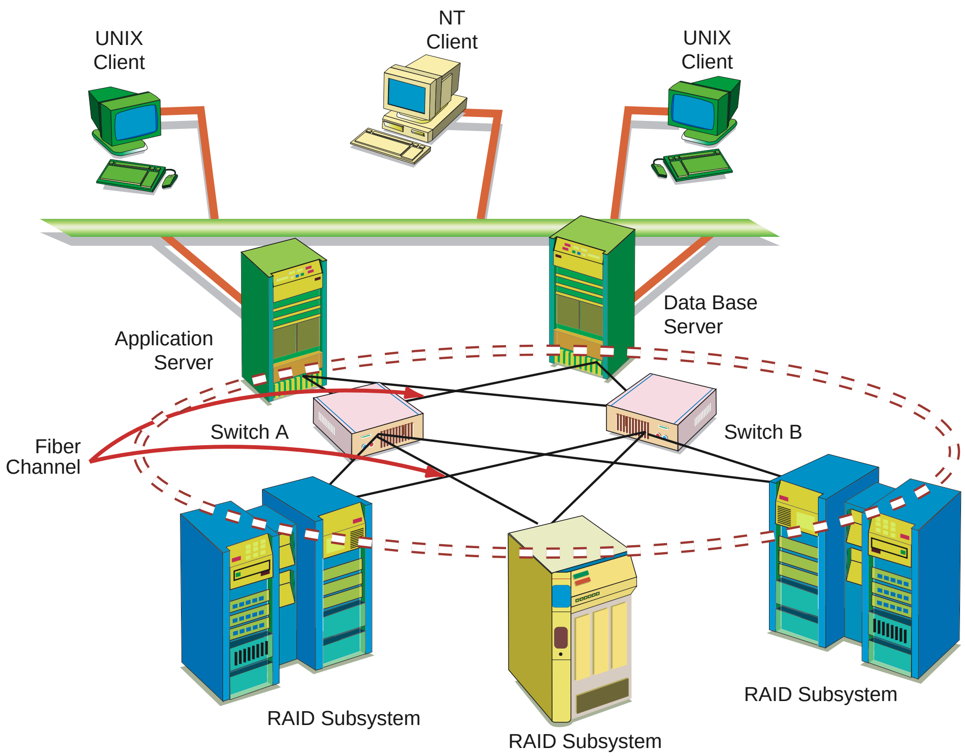

Figure 11 - A heterogeneous Fibre Channel Storage Area Network or FC-SAN

Challenges in managing FC-SANs

Now that the complexity and scope of the Fibre Channel SAN environment is better understood, the challenges of managing that environment can be considered. The basic challenge stems from the fact that Fibre Channel SANs utilize a complex network for interconnecting storage devices and server systems. This network is potentially made up of multiple components that have both physical and logical relationships to one another. A breakdown in one component or link in the SAN may manifest itself differently depending on the component that recognizes the condition.

Problem isolation is extremely difficult with FC-SANs.

However, since a SAN is a network, there are many benefits from deployment in the area of systems and network management. From a fault management perspective, tools and mechanisms can control data from multiple SAN components from a central administration node. Without the means to correlate data in a SAN, problem isolation is extremely difficult with RFC-SANs, if not impossible. For example, if multiple link failures were to simultaneously occur in a Fibre Channel network, these would need to be correlated with the multiple I/O failures seen in processing a host's I/O requests. New IP probably the greatest challenge involves the configuration management of SANs. Due to the large number of components, and the multitude of physical and logical relationships to one another, robust configuration capabilities must be provided for the Fibre Channel SAN.

There are tradeoffs between keeping track of configuration information within the Fibre Channel network, or forcing a central management platform to support Fibre Channel network topology using management mechanisms such as configuration files, name services, and Simple Network Management Protocol (SNMP) interactions.

There are also challenges associated with performance management of FC-SANs. Performance information must be provided at a component level as well as at the overall system level. Tools and capabilities must exist that again correlate data from a variety of components to provide a system level view of the overall SAN performance. Common capabilities must be provided to allow software and firmware updates to be managed from a central management station. These capabilities should allow a generic mechanism to transport updates to components from different vendors.

For the support of accounting or asset management, capabilities such as standardized SAN resource identifiers containing asset information must be defined. In addition, a common mechanism must be provided to obtain the asset information from SAN resources. Cluster management is an inherent requirement of SAN management and imposes its own set of challenges. Managing SAN clusters calls for robust solutions for workload management and for single point of control of the cluster's component systems. Data and storage management must be supported within and across clusters.

Security management in the FC-SAN also creates challenges.

Creating a scalable, single operational management view for SAN clusters enhances productivity and application availability by enabling the automatic detection, monitoring, and control of key cluster components. Security management in the SAN also creates challenges. For example, restrictive mechanisms such as zoning can have both beneficial and negative consequences for central management. While zoning provides a means to isolate certain storage devices, a centralized manager may need to be a member of all zones and must therefore be a "trusted" entity.

The overall security of the SAN then becomes only as secure as the trusted entity. Ultimately the purpose of providing robust management capabilities for the Fibre Channel SAN environment is to support policy management. Policy management is becoming increasingly important for the enterprise where Fibre Channel SANs will most likely be implemented. In the SAN environment, management capabilities must support the invocation of services chained to well defined policy. Examples of services that require policy are:

- Disaster recovery
- Back-up
- Hierarchical Storage Management (HSM)
- Archival

Benefits of FC-SANs

The benefits of FC-SAN are shown in Table 8. They include the important features of centralized data management and LAN and server free backup. Since backup windows are the single largest problem reported in a recent ITCentrix survey of storage managers, SAN provides a benefit that directly addresses this important item.

Benefits of a FC-SAN Storage Architecture	
Centralized management of data	Yes
Remote connectivity of storage	10 kilometers
Reduced storage management costs through consolidation	Yes
Ease of scaling (adding) storage capacity and processor performance	Yes
Separation of storage and server purchases	Yes
Centralization of storage management and support	Yes
Avoidance of single server bottlenecks	Yes
Avoidance of networking bottlenecks	Yes
Removal of backup data flow from the LAN	Yes
Server-free backup	Yes
Increased storage resource utilization	Yes
Improved data availability	Yes

Table 8 - The benefits of FC-SAN

FC over IP Industry initiatives

To try and offset the 10-kilometer distance limitation of Fibre Channel, switch vendor Gadzoox Networks Inc. of San Jose, California and Lucent Technologies Inc., of Murray Hill, NJ, are jointly developing the IPFC specification (IP Fibre Channel). IPFC is a specification for the routing of Fibre Channel traffic (requests and data) over IP networks to connect islands of SANs separated by long distances. This is shown in Figure 12.

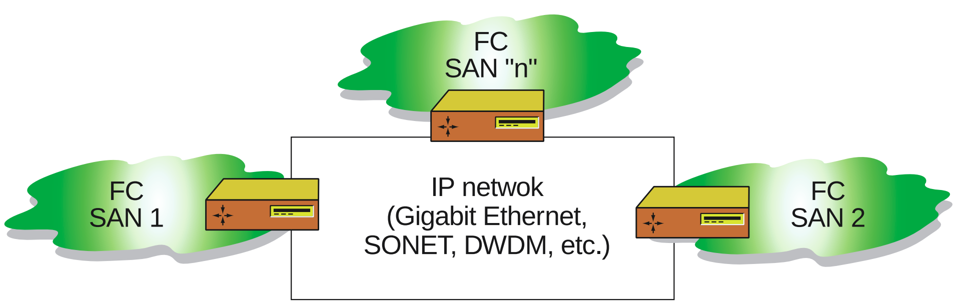

Figure 12 - Specialized routers link Fibre Channel SAN "islands" over long-haul IP networks such as WANs or MANs

The proposal involves sending encapsulated FC packets over IP networks as shown in Figure 13.

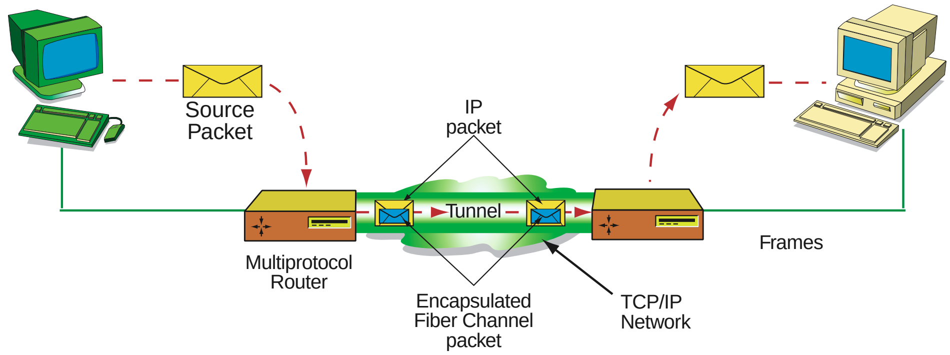

Figure 13 - An encapsulated FC packet over an IP network

Cisco and Brocade's plan to encapsulate FC over IP

Another effort to overcome Fibre Channel distance limitations is evident in the partnering of Cisco Systems. with Fibre Channel switch manufacturer Brocade Communications Systems Inc., both of San Jose, California, to encapsulate the Fibre Channel protocol in IP, in order to join together SAN islands using standard networking. As reported in the June issue of InfoStor, the companies anticipate products in the second quarter of 2001.

Ethernet (What is it?) and Ethernet Storage Area Networks (E-SAN)

7

The different types of Ethernet defined

The investment can be leveraged by using E-SAN.

Table 9 lists the different varieties of Ethernet. The first number in the name of the variety indicates the speed in Mbits/sec and the last number refers to the meters per segment (multiplied by 100). *Base* stands for baseband and *Broad* stands for broadband. 10Base-2, 10-BaseT, 100BaseT and 1000BaseT (Gigabit Ethernet) and 10-Gigabit Ethernet are extensions of the IEEE 802.3 Ethernet networking standard. In general the big advantage of Ethernet is that it is ubiquitous. Most organizations already have an Ethernet infrastructure in place. The investment in this infrastructure can be leveraged to provide a storage infrastructure by using E-SAN. This is a major feature additional advantage over to Fibre Channel, which nearly always requires significant investment in equipment, interfaces, network analyzer tools and training.

Table 9 - The different varieties of Ethernet

Different varieties of Ethernet	**Type of Cable**	**Maximum segment length**
10Base-5	Thicknet coaxial cable	500 meters
10Base-2	Thinnet coaxial cable (RG-58 A/U	185 meters
10Base-T	Twisted-pair cable	100 meters
10Base-F	Fibre optic cable	4kilometer backbones
100Base-TX	Two pairs of cable category 5 UTP or Category 1 twisted pair	100 meters
100Base-T4	Four pairs of Category 3, 4 or 5 UTP wiring	100 meters
100Base-FX 100VG-Any LAN	Fibre Optic cable Hierarchical-twisted pair	
1000Base-T (Gigabit Ethernet)	Category 5 copper cable	100 meters
1000Base-LX (Gigabit Ethernet)	Multimode fiber-optic cable Single mode fibre-optic cable	550 meters = multimode 3000 meters - single mode
1000Base-SX (Gigabit Ethernet)	Multimode fiber-optic cable	300 meters = 62.5 micron multimode 550 meters = 50 micron multimode
1000Base-CX (Gigabit Ethernet)	Twisted pair copper cable	25 meters
10,000Base-T (Ten Gig Ethernet)	Fibre optic and copper cable	Various depending on variety

Part of the reason for the wide acceptance of Ethernet has been the smooth transition between different generations. To ease the transition from 10base-T to 100base-T, most network interface cards (NICs) support both speeds and negotiate to pick the fastest speed supported by both endpoints of a connection. Similarly, the introduction of 1000base-T is being accelerated by autoranging 100/1000 and 10/100/1000 NICs and switches.

Gigabit Ethernet functional layers

Media Access Controls or MACs are the rules defined within a specific network type that determines how each station accesses the network cable. Using a token-passing method, a carrier sensing and collision detection method or a demand priority method prevents or detects simultaneous access to the cable. The MAC used for 100BaseT Ethernet is based on the "demand priority" access method in which the central hub scans all its ports in a round-robin fashion to detect stations that want to transmit a frame. Higher priorities can be requested by ports to transmit real-time information like video or audio. ThMAC sublayer is shown in Figure 8.

Figure 14 – The MAC sublayer of Gigabit Ethernet

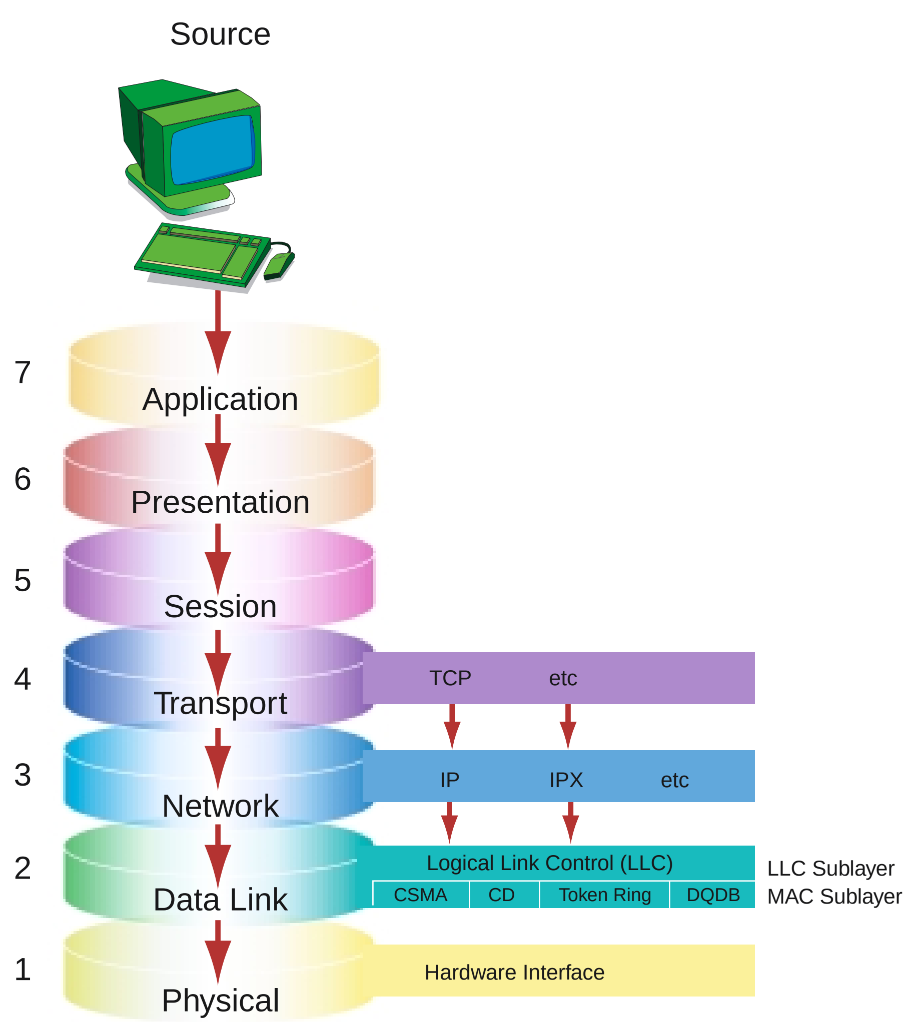

Configurations of E-SANs

There are five types in E-SAN configurations. These are listed below and illustrated on the cover of this report and the Figures 15-17. It is important to recognize that E-SAN offers more flexibility in network configurations than FC-SAN, which currently is effectively limited to homogeneous (and proprietary) short distance (10 kilometer) deployments. Conversely E-SANs can be heterogeneous, short or long haul networks, and with a mix of Fibre Channel and Ethernet protocols.

	E-SAN	FC-SAN
Campus configuration with Separate LAN and SAN	✔	✔
Campus configuration with combined LAN and SAN	✔	
Homogeneous configurations	✔	✔
Heterogeneous configurations (switches, storage)	✔	
WAN configurations in remote geographies	✔	

E-SAN configuration with separate LAN and SAN

E-SANs can also be configured whenever there is a separate Ethernet network for block storage traffic as shown in Figure 15. The SNMP bridge is needed only if there is a desire to manage the SAN from a machine not directly connected to one of the SAN switches.

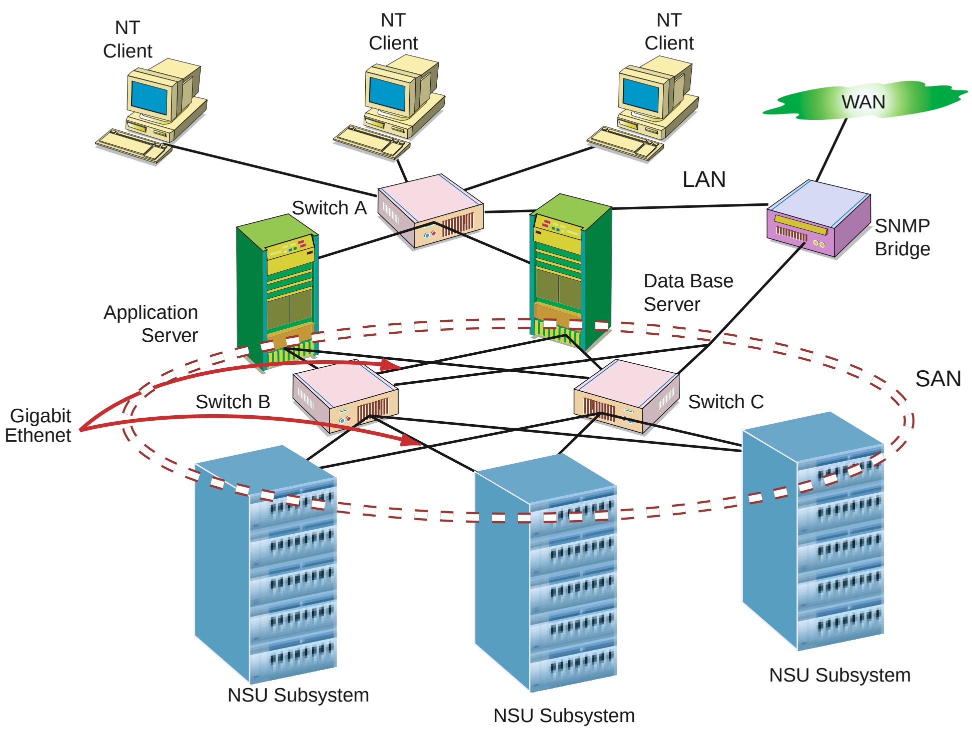

Figure 15 – A homogeneous E-SAN with separate LAN and SAN

E-SAN configuration with combined LAN and SAN

Another option is to configure the E-SAN in a way that block storage traffic shares a single Ethernet network with other LAN message traffic as shown in Figure 16. In either case bandwidth must be matched to the configuration and the network workload and performance objectives. A major advantage of the E-SAN flexibility of configuration is that the benefits of SAN can be delivered to smaller installations (e.g., remote offices) where there is only one network, without having to invest in separate equipment for a dedicated E-SAN to separate block traffic from LAN message traffic.

Figure 16 – A homogeneous E-SAN with combined LAN and SAN

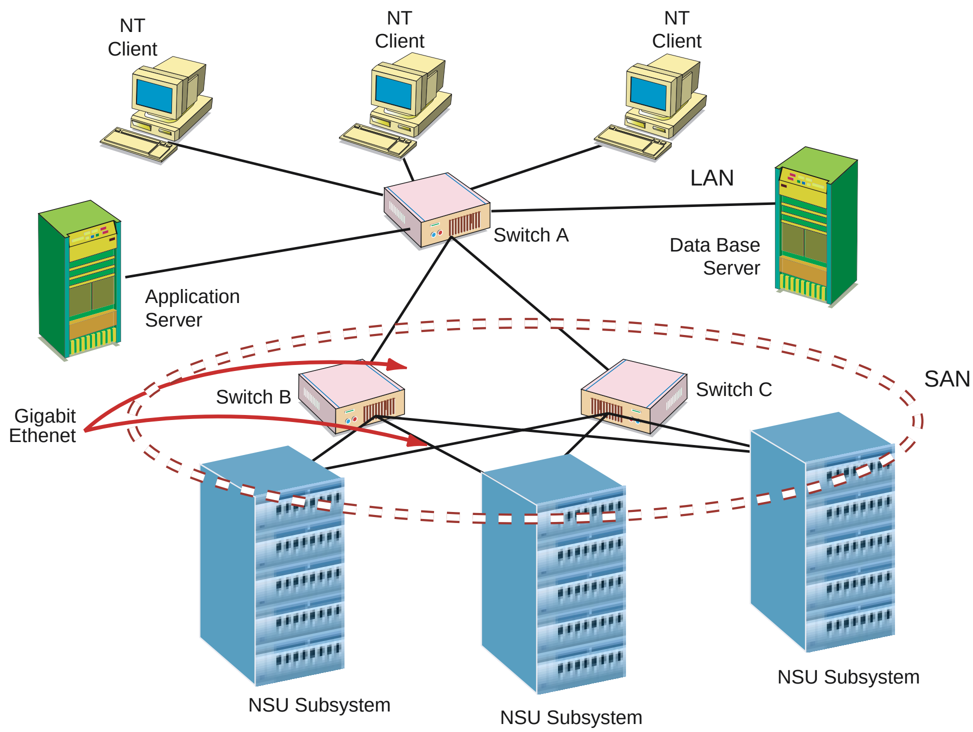

Interconnected remote E-SANs in different cities

Perhaps the biggest advantage of E-SAN is it's long distance capability

Perhaps the biggest advantage of E-SAN is the capability of the Internet Protocol (IP) for long distance networking. The increasing need to push data across a WAN has led some vendors and their customers to explore what they think is a better idea than FC-SANs using the Internet Protocol to manage storage. The Internet explosion has increased the need for companies to manage huge amounts of data, not just within an enterprise but across a series of broadly distributed sites. This is among the major needs addressed by 3Ware's NSU™. "Storage over IP could be a tenfold explosion in terms of how storage can be connected. It will push Fibre Channel into a niche in the data center," said Michael Brown, Chairman and CEO of Quantum Corp., a Milpitas, Calif., storage device maker.

An example of how remotely connected E-SANs would benefit an enterprise is when there is a need to run remote backup of data to a remote geography. This and other suitable applications for the 3ware NSU are discussed in Chapter 9.

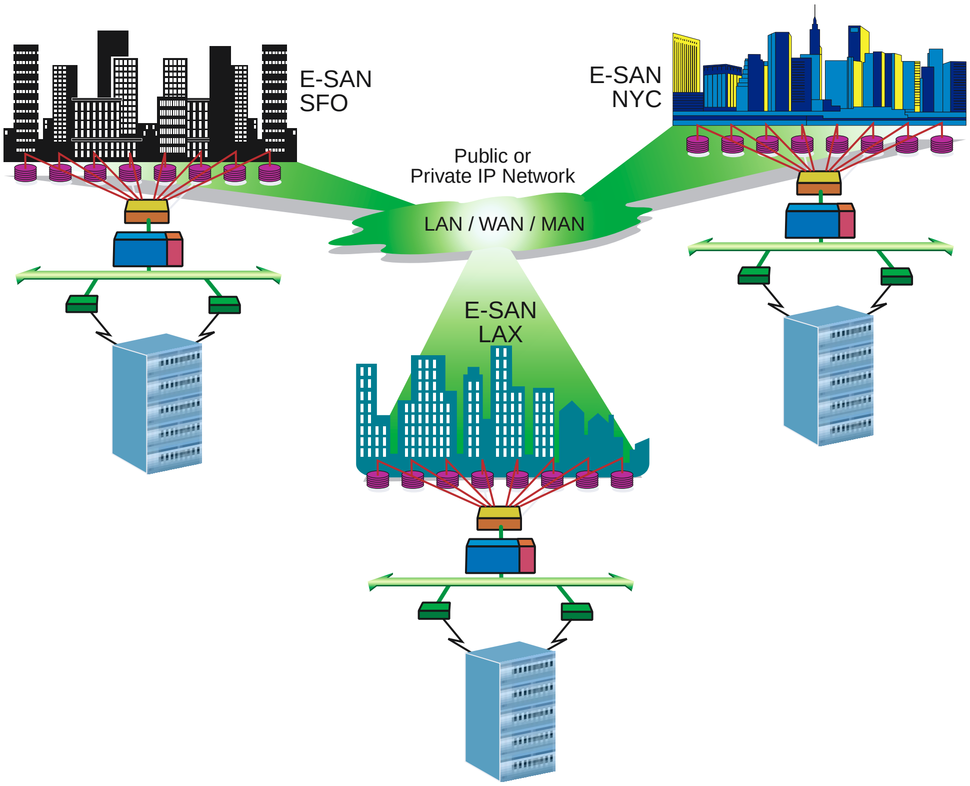

Figure 17 – Interconnected remote E-SANs in different cities

Heterogeneous Ethernet Storage Area Networks or E-SANs

Unlike FC-SANs, E-SANs can be easily configured with either homogeneous or heterogeneous components. This is because Ethernet switches interoperate with each other using well-understood standards for security and congestion control.

Figure 18 - A heterogeneous Ethernet Storage Area Network or E-SAN

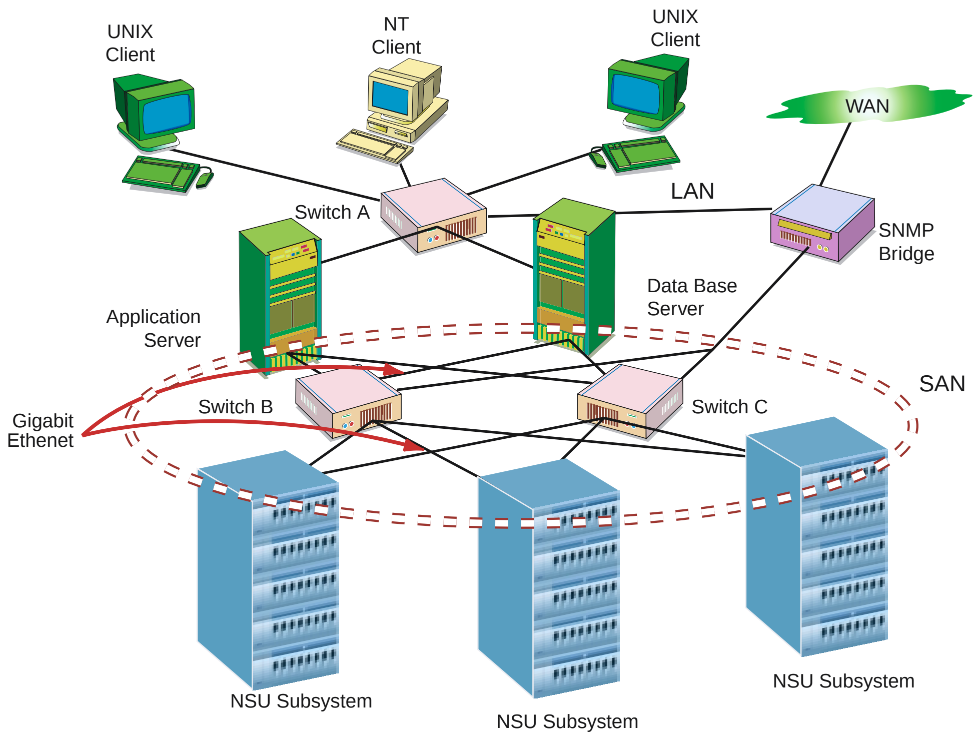

Overview of SCSI over IP and the different approaches to performance

IP was designed to handle high volume traffic with nonsequential data packets that need to be reordered at their destination. SCSI, however, needs to have packets arrive in order. If packets don't arrive in order, a lengthy higher-level retry must be invoked for the entire operation. A frequently encountered criticism of SCSI over IP therefore is the notion that performance will suffer as packets are reassembled for execution at the SCSI storage system. 3ware tackles this problem by using a new command set and patented internal packet switched (Disk Switch™) architecture described in Chapter 9. There are other approaches to this problem, as SoIP, STP, and iSCSI are all trying to solve the performance issues of SCSI over IP.

The iSCSI initiative of Cisco, IBM, HP, and others

Cisco Systems Inc., of San Jose, Calif., and IBM, of Armonk, N.Y., have teamed up to address a standard for flowing SCSI commands over IP. The companies recently delivered a draft specification for the so-called iSCSI technology to the IP Storage working group, a group within the Internet Engineering Task Force (IETF) standards organization. Recently, Cisco announced plans to acquire NuSpeed Internet Systems Inc., a Maple Grove, Minn., company that is developing iSCSI products.

The iSCSI specification assumes the presence of NICs that include hardware to accelerate the protocol. While this is likely to increase performance, it also means that the specification must be extremely stable before manufacturers can commit to the design of the new integrated circuits needed for the NICs. The cycle from initial specification to consensus, ratification and working parts may take several years.

3ware has taken a different approach and has implemented its SCP protocol in software using standard NICs. This approach provides an immediate solution with a working protocol for customers that need alternatives to DAS or FC-SAN today, yet allows for software upgrades to iSCSI when the standard is completed and implemented.

3ware has implemented its SCP protocol in software using standard NICs.

The SoIP initiative of Nishan flows SCSI block commands over UDP/IP

Another potential stumbling block for using IP to deliver storage data is performance. Processing the IP stack in software alone puts a high demand on a server's CPU cycles.

To solve that problem, startup Nishan Systems is taking a different approach. The San Jose company announced plans to ship before 2001 an end-to-end storage over-IP solution that uses IP both in the WAN and inside the data center. It will accomplish this by embedding the IP stack into a network adapter thus bypassing the server's CPU.

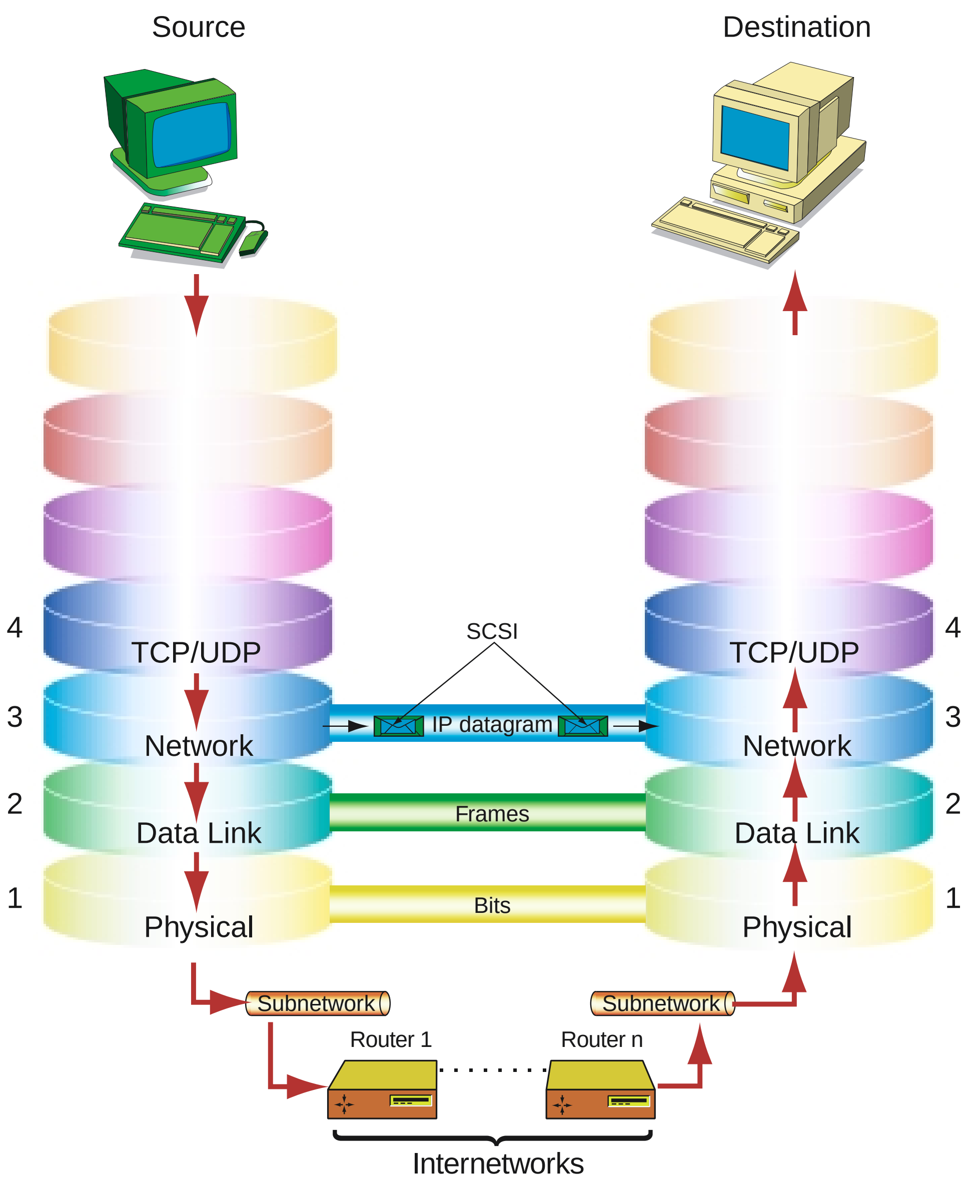

Figure 19 – Nishan's SoIP (SCSI over IP) proposes to run SCSI block commands over UDP/IP

Adaptec's EtherStorage initiative

Adaptec Inc., of Milpitas, California has spent two years developing its new EtherStorage technology and has also submitted to the IETF a draft specification called SEP (SCSI Encapsulated Protocol), that enables block-based storage traffic to be transferred over existing IP and Ethernet networks.

Figure 20 - Adaptec's SAN Transport Protocol (STP) would replace TCP in the stack for SCSI over IP

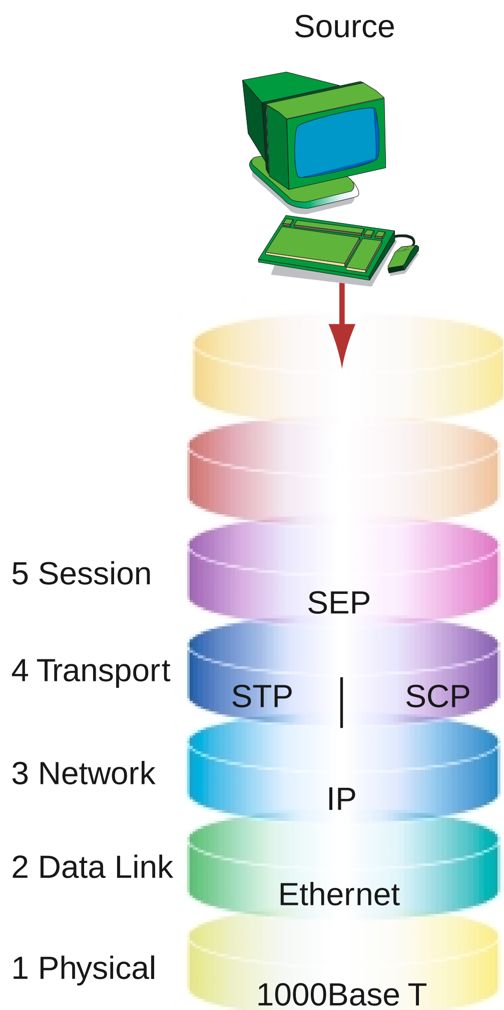

A taxonomy of SAN

In order to classify the various SAN initiatives, 3ware's VP of Research and Technology, Robert Horst has proposed a taxonomy for clarity. This taxonomy is shown in Table 10. It is important to note that only with 3ware is one board required for the NICs in the server and not multiple boards, as is the case with all other SAN levels.

SAN Level	Description	Switch	Network NIC	Storage NIC	Examples
0	Legacy - two networks	Ethernet and Fibre Channel or bus	Ethernet	Fibre Channel or SCSI	FC-SAN vendors such as EMC, Sun, HP, Compaq
1	Multiprotocol Switch	Combined Ethernet and Fibre Channel	Ethernet	Fibre Channel	Nishan and Cisco (NuSpeed)
2	FC over IP	Gateway from Fibre Channel to Ethernet	Ethernet	Fibre Channel	SAN Valley, Lucent, Entrada Networks
3	SCSI tunneling over IP	Ethernet	Ethernet	iSCSI	iSCSI working group
4	Native IP Storage	Ethernet	Ethernet		3ware's NSU

Table 10 - A taxonomy of SAN

The Advantages of E-SANs over FC-SANs

8

The Promises versus the Realities of FC-SAN

When first initiated by the storage industry, FC-SAN had many lofty objectives. The concept of having a dedicated private network (optimized for storage traffic) that separates LAN and storage traffic is an excellent one. However this initiative has not developed to the point where an agreed to set of standards has emerged. As a result, there is currently no interoperability between SAN implementations, which has created a set of proprietary FC-SAN solutions.

Unlike 3ware, most storage vendors are not experts in networking technology.

Storage vendors are not experts in networking technology and there were many aspects of networking that were left to implementation by individual vendors. Two major examples of this are security and congestion control, which are still not standardized for FC-SAN and may not be done so for years.

The vision of FC-SAN, therefore is more vision today than reality as shown in Table 10.

Table 11 - The promises versus the realities of FC-SAN

Fibre Channel Storage Area Network (SAN)	
Promises of FC-SAN	**Realities of FC-SAN**
• Easier centralized management • LAN free back-up • Higher application availability • Block level transfers • Better application performance • Resource sharing • Unlimited Scaling	• Requires new network infrastructure = $$$$ • Lack of administrators • Additional training = $$$$ • Network analyzers = $$$$ • Standards two years away • Interoperability "rats nest" • Slow adoption • Frequent buyer's remorse

Deterrents to FC-SAN today as reported in a recent survey in Computerworld

Figure 21 illustrates data published recently in an article in Computerworld 8/7/00 by Kathleen Ohlson and Sami Lais titled "Skittish Users Say Hype About SANs is 'Baloney.'" The article talks about the many problems of FC-SANs today as reported by early SAN aware professional IT managers.

It is critical to understand that E-SANs while still in development, will be free from many of the problems shown in Figure 21 and discussed below.

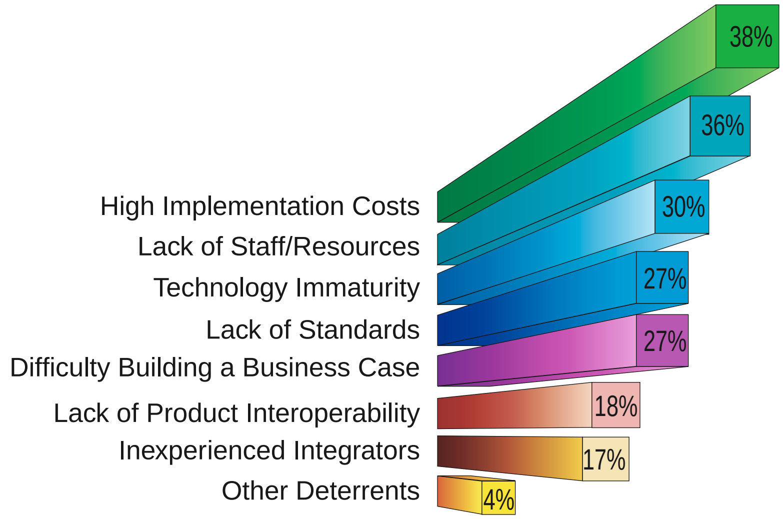

Figure 21 – Results of a survey regarding deterrents to FC-SAN deployment as reported in August 7, 2000 Computerworld

Why FC-SAN deterrents won't apply to E-SAN

Table 11 presents the deterrents reported by SAN implementers in the August 7, 2000 edition of Computerworld and explains why these deterrents will not apply to E-SANs. Examining this table shows that the complaints about FC-SAN can be grouped into three main categories, 1) High costs, 2) lack of standards and 3) the learning curve of a new technology. Since the investment required for E-SAN is dramatically less, and Ethernet is a robust standard well known by all network administrators, it is very clear that E-SAN is not FC-SAN in that it avoids all these common objections.

Consideration of SAN Deployment[2]	**Deterrent to FC-SAN**	**Deterrent to E-SAN**
High implementation Costs (38%)	Yes- Requires entire new Fibre Channel fabric **Can be big $$$.**	Much less - Uses existing network technology. May require upgrade.
Lack of staff / resources (36%)	Yes- plus training and capital investment in network analyzers **Can be big $$$.**	Much less since Ethernet networks are standard and widely deployed.
Technology immaturity (30%	Yes - arguments over standards	In some areas - not others.
Lack of) standards (27%)	Yes - Especially in security, congestion control and switches.	No - Well established security, congestion control and switch specs.
Difficult cost justification (27%)	Yes - High costs raise the bar.	No - Lower bar due to lower costs to implement.
Lack of product interoperability (18%)	Yes - Very expensive to implement and test. **Can be big $$$.**	Not at the switch level where most FC-SAN problems are today.
Inexperienced integrators (17%)	Yes	No - Ethernet is well understood. As is SCSI driver installation.

Table 12 - Why FC-SAN deterrents won't apply to E-SAN

[2] Refers to the survey published in the August 7, 2000 edition of Computerworld.

3ware's Network Storage Unit (NSU)

9

Gigabit Ethernet Storage has arrived

The 3ware Network Storage Unit (NSU) is the first storage solution to support storage over IP. It provides an all-new approach to implementing a SAN solution. With the 3ware NSU, businesses can now leverage their existing investment in Ethernet network infrastructure, management software and people to build and maintain an enterprise level storage solution.

With the 3ware NSU, businesses can now leverage their existing investment in Ethernet network.

With the 3ware NSU there is no need to create an all-new network requiring a major investment in cabling, switches hubs and tools in order to realize the advantages of SAN. The 3ware NSU fits into existing networks, providing the cost savings and manageability of virtual consolidation where multiple servers on a network can share a common storage resource.

Ease of Implementation

The 3ware NSU is a modular storage unit that is scaleable from the work group to the enterprise. It can be seamlessly integrated into an existing LAN or configured on a separate subnet using standard CAT5 cabling. Storage appears as SCSI to the host so existing SAN management tools can be employed. A low cost Ethernet port of either 10/100 or 1000baseT is the only hardware that is required to support the installation on the host side. 3ware will provide a client side driver for Windows 2000, Windows NT or Linux 2.2. Additional client drivers for Solaris, MacOS, and Linux 2.4 are in development.

Enterprise Class Data Protection

The 3ware NSU has been designed to provide an enterprise with the highest level of protection for data. The patented 3ware DiskSwitch and TwinStor technologies are supplemented by redundant power supplies, RAID 1, 5 and 10 support, hot swappable components, and hot spare drive features. The use of multiple Ethernet connections can provide alternate path support should a cable connection be interrupted. It is possible with 3ware technology to have individual switches connected to NSUs in a network connection, and then aggregate these switches in redundant configurations for optimal data protection.

The 3ware NSU has been designed to provide an enterprise with the highest level of protection for data.

Ease of Network Management

The 3ware NSU is compatible with standard network management systems and can be deployed using the same tools currently in use. With the NSU there is no need for additional training or the purchase of new networking equipment and software as would be the case with FC-SAN.

High Performance

The 3ware NSU provides performance equal to that found in today's distributed computing environment. The 3ware NSU can be configured as a direct attached device or in an E-SAN. As with most applications, data requirements such as block size and availability of network bandwidth are the controlling factors to overall data throughput. With the 3ware NSU performance equal to or greater than Ultra SCSI can be expected but at a much lower cost than Fibre Channel.

Connectivity and Interoperability

Since the 3ware NSU utilizes Ethernet, the distances between storage connections are no longer an issue. Industry standard CAT5 cabling can be used to provide 100 meters distances between most arrays and switches and additional 100 meters to any server. Along with standard fiber optics, that distance can be extended to 8 Kilometers between switch points. Given that the 3ware NSU capitalizes on the use of Ethernet (which has a defined set of standards), customers are free to choose the switch of their choice in supporting a network.

Infinite Scalability for Unlimited Capacity

The 3ware NSU solution can scale virtually infinitely to meet any business's increasing demand for storage. Each rack mounted or tower NSU currently has an unformatted capacity of either:

- 240 GB with 8 x 30 GB drives, or
- 600 GB with 8 x 75 GB drives

Whether rackmount or tower, each NSU unit contains locking disk bays, hot swap drives and power supplies and hot spare. Given that the 3ware NSU operates as a switched node on an IP network, the only limitation on the scalability of the 3ware NSU solution is the network infrastructure switching capacity. A 12 port switch for example can provide scaling capacity to 6.6TB (600GB X11).

Network and system management tools

3DM provides full function GUI based management.

3ware has created its own 3DM (3ware Disk Manager) software tool to remotely monitor and manage the health of the NSU through a full function browser-based GUI. With 3DM a storage array can be managed remotely. Standard SNMP compatible MIBs are supported by the 3ware NSU - allowing an easy integration into your existing framework management infrastructure.

Total Cost of Ownership

Any business considering a SAN implementation needs to look at the bottom line. It is not only the cost of the array but also the networking infrastructure to support a SAN. Unlike other SAN vendors, 3ware understand not only storage but also networking. Therefore, in developing the NSU, 3ware decided to utilize the existing Ethernet network that is well understood by existing network administrators and does not require learning new skills specifically reserved for storage administrators. The bottom line impact of using Ethernet allows the user to extend the benefits of SAN to the department using an existing LAN. It is no longer necessary to purchase a new set of tools and procedures in order to gain the benefits of SAN as is the case with FC-SAN.

Platforms ideal for connecting an NSU

- High performance server platforms
- Platforms that requires SAN functionality.

The NSU's patented internal packet switched architecture

The 3ware NSU "looks like" a regular SCSI controller to supported hosts running Windows NT, Windows 2000 or the Linux operating systems. When the NSU is installed, a driver is installed in the host that sends customized 3ware storage commands (using the SCP protocol) over Gigabit Ethernet to the NSU. This command set[3] is a subset of the SCSI command set and provides both simplicity and increased performance compared to SCSI. The NSU recognizes the SCP command and converts it to an ATA protocol operation for transmission to the backend ATA drives over their private 100MByte/sec busses. RAID 0, 1, 5, 10 and JBOD are supported for a broad choice in availability and performance.

A driver is installed in the host and sends customized 3ware storage commands over Gigabit Ethernet to the NSU.

ATA/IDE drives enable 3ware to use a multiplexing data packet switching technique not available with SCSI. This technology is patented and known as DiskSwitch. Disk Switch allows for the simultaneous data streaming from multiple drives in an array. The 3ware architecture is a truly independent I/O bus architecture that eliminates SCSI bus arbitration delays since all drives in the array stream data simultaneously on private high-speed busses. The internal NSU architecture is shown in Figure 22.

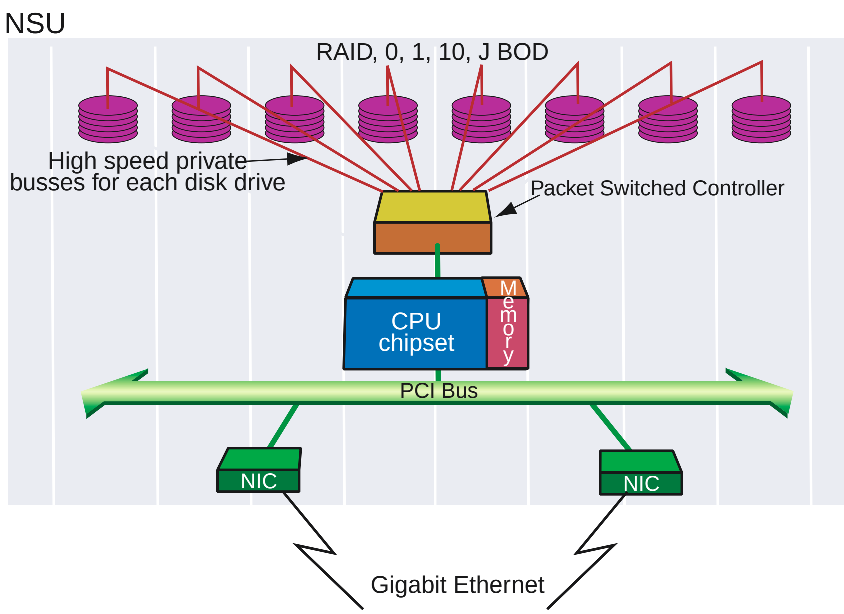

Figure 22 - The 3ware NSU is based on a patented internally switched architecture

[3] The 3ware command set is part of the Storage Control Protocol (SCP) that is a 3ware developed storage protocol for sending data requests and data blocks over TCP/IP networks.

The NSU is available in Rackmount or Tower configurations

The 3ware NSU is available in both rackmount and tower configurations as shown in Figure 23.

Figure 23 - The NSU is available in Rackmount or Tower configurations

Availability and RAID configurations

The NSU supports a full range of RAID configurations. Each NSU can support RAID 0 (striping), RAID 1 (mirroring), RAID 5 (parity RAID) and RAID 10 (striping and mirroring) as well as IBOD. High-speed non-volatile disk buffer memory accelerates RAID functions, particularly disk writes. These configurations are shown in Figure 24.

Figure 24 - The NSU is available in RAID 0, 1, 5 and 10

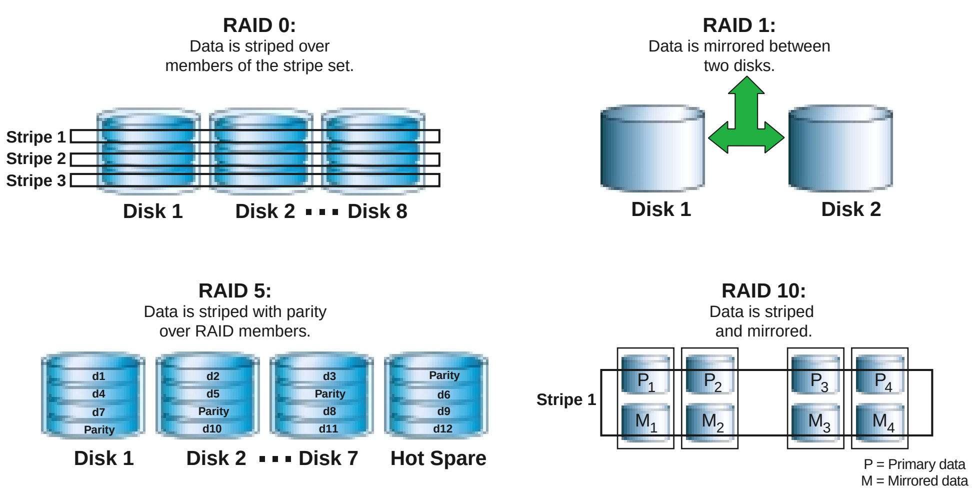

RAID 0 – Striping for optimal performance

RAID 0 accelerates disk access by striping data across the disk array, thereby spreading I/O evenly across multiple disk spindles. RAID 0 provides the best overall performance but provides no resilience to disk errors or failures.

RAID 1 – Mirroring for optimal availability

RAID 1 arrays consist of pairs of disks in which each disk is maintained as an exact replica or mirror of the other. This method provides the maximum data redundancy and resilience. It is also the most expensive RAID alternative since it requires disk space equivalent to twice the usable capacity. Performance of RAID1 is generally very good for random I/O, because both disk arms can simultaneously seek to different records. For sequential I/O, RAID1 is normally similar to a single non-mirrored disk. However, 3ware has a technique, called TwinStor, that greatly improves sequential read performance of mirrored disks.

RAID 5 – Distributed Parity for cost effective data protection

In RAID 5 arrays, data is striped across arrays in a fashion similar to RAID 1, but RAID 5 provides fault resilience by keeping parity information on each stripe of data. If a failure occurs, the contents of that block can be recreated by reading back the other blocks in the stripe along with the parity. Parity information is distributed throughout the array to minimize potential bottlenecks. The capacity overhead of RAID 5 is equivalent to one disk drive regardless of the size of the array. RAID 5 offers low cost redundancy, but has lower performance than RAID 1. RAID 5 is often configured with an additional spare drive. The spare helps to reduce the time the array is exposed to double failures, and also reduces the time the array is running with degraded performance.

Parity information is distributed throughout the array.

RAID 10 – For optimal availability and performance

RAID 10 configurations offer the advantages of both RAID 0 (striping) and RAID 1 (mirroring) as shown in Figure 24. Although more expensive to deploy because of the mirroring of each disk unit, RAID 10 configurations provide the best of both performance and availability of data.

RAID – Hot Sparing

In the event of a disk failure, RAID 1, 10 or 5 arrays a hot spare can be rapidly and automatically built using available "hot-spare" drives.

Network Connections

- Standard TCP/IP Protocol
- 10/100/1000 Mbits /sec
- CAT5 or fibre optic cabling

Performance and Scalability

- Available in 240 or 600 GB modules
- Scalable to outboard switch capacity
- Exceeds external Ultra SCSI performance
- Enterprise level data availability
- Hardware RAID 0,1,5 and 10 support
- Hot swap drives and power supplies
- Hot spare drive support

Ease of Management with 3ware's 3DM software utility

- SNMP compatible MIB support
- Email notification of critical events
- Network based management tool
- Browser based
- 3ware NSU monitoring and management
- Remote configurability
- Secure management interface to prevent unauthorized access
- Packaging and support
- Tower or rack mount available

Operating system support

- Windows NT
- Windows 2000
- Linux 2.2
- Solaris, MacOS, FreeBSD, Linux 2.4 all in development

E-SAN and DAS Configurations of the NSU

The NSU can be configured as in either SAN or DAS.

The NSU can be configured in either SAN or DAS configurations as discussed above. These various configurations are shown in Figures 26-29. The NSU can be configured in a Gigabit Ethernet E-SAN with an E-SAN separate from the LAN (Figure 26), or in an E-SAN configuration with the SAN and LAN combined (Figure 27). Importantly the product can also be configured as Direct Attached Storage in configurations where the server has no SCSI slots [and has free NIC slots] or in configurations where the cost advantages of the NSU are desired (Figure 28). Remote NSU based E-SANs can be connected using WAN technology as shown in Figure 29[4].

[4] Remote SAN connection is not possible with FC-SANs due to the 10 kilometer limitation of Fibre Channel. See also the discussion in Chapter 6 regarding tunneling SANs and Figure 12.

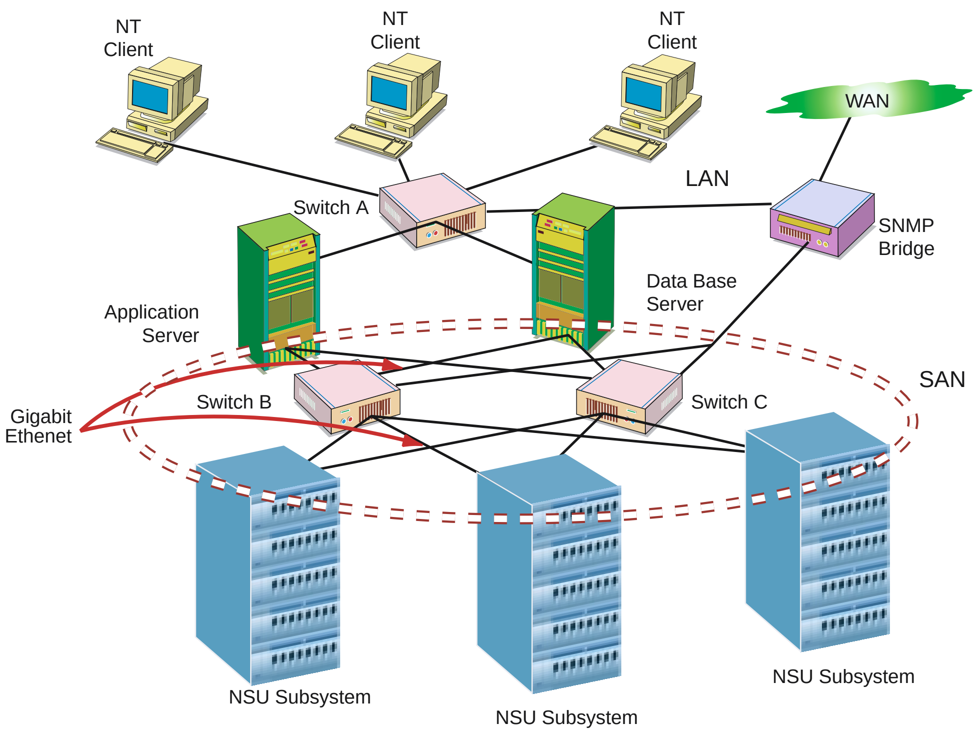

Figure 25 – The NSU configured in a Gigabit Ethernet E-SAN with an E-SAN separate from the LAN

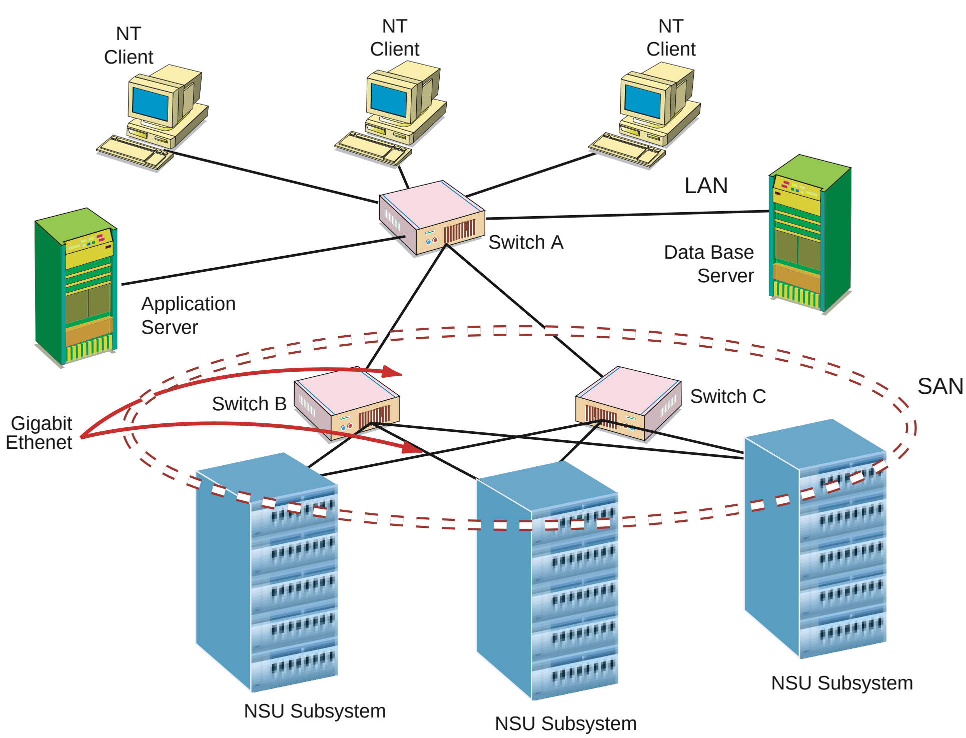

Figure 26 – The NSU configured in a Gigabit Ethernet E-SAN with an E-SAN and LAN combined

External Gigabit Ethernet connected DAS

The NSU provides SCSI performance but with cost effective ATA drives instead of more expensive SCSI drives.

Compared to DAS SCSI storage, the internally and externally switched architecture of the NSU provides SCSI performance but with cost effective ATA drives instead of more expensive SCSI drives. In DAS configurations, the NSU is compatible with existing networking infrastructure, and is highly scaleable. The NSU in E-SAN configurations easily exceeds distance limitations of differential SCSI. A DAS configuration for the 3ware NSU is shown in Figure 27.

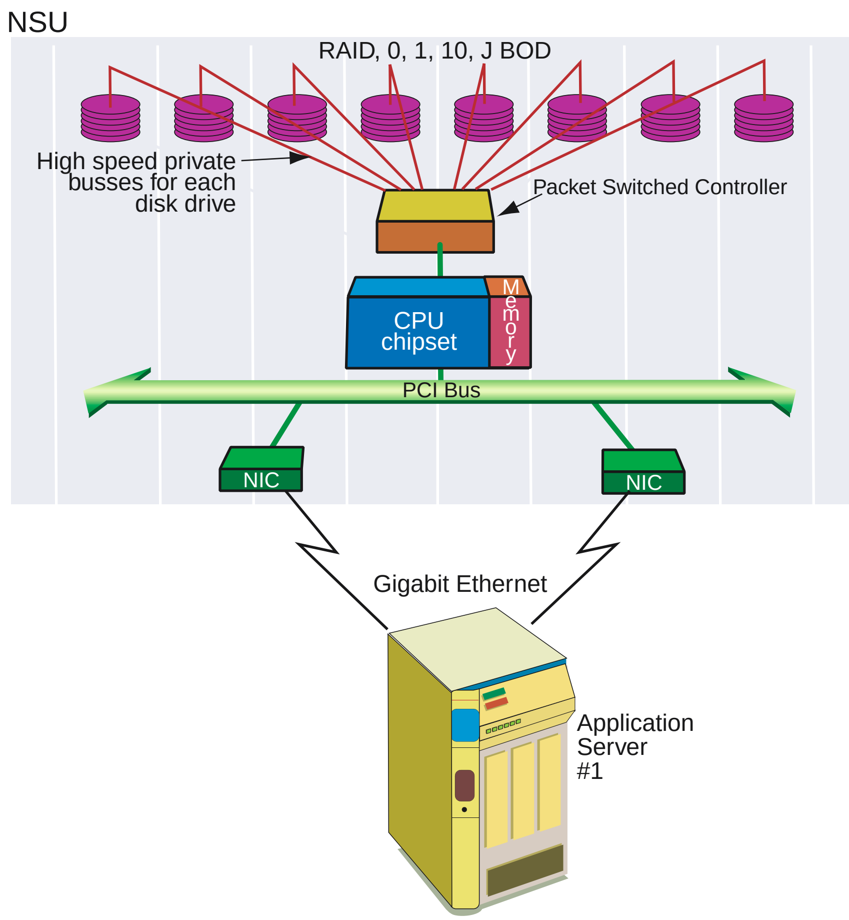

Figure 27 – The NSU configured as Direct Attached Storage (DAS) with redundant NICs

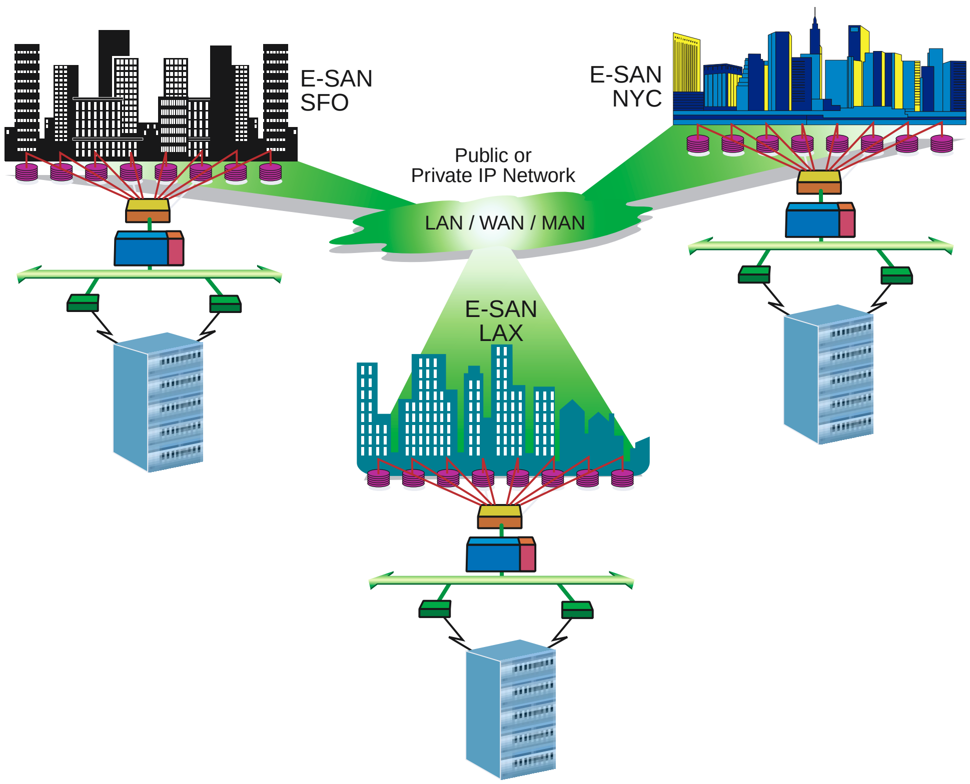

Figure 28 – Interconnecting Remote NSU based E-SANs

ATA drives give the NSU unmatched storage Cost/Performance advantages

Driven by the demand for desktop PC storage, ATA/IDE drives have become the highest unit volume disk drives as shown in Figure 30. This results in an unprecedented economy of scale that makes them the most cost-effective storage component available. Because of defacto PC standards, all of these drives share a high degree of compatibility and standardization across many vendors. High volume manufacturing processes, coupled with standardization, leads to greater consistency and quality in these drives, as well as lower cost.

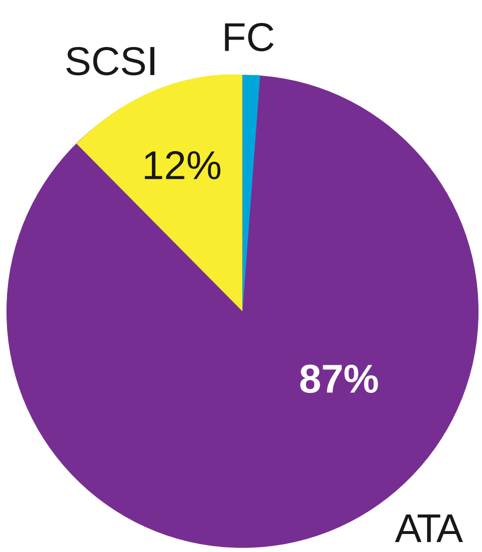

Figure 29 – High volume ATA production results in the lowest possible drive costs

Ethernet gives the NSU unmatched network Cost/Performance advantages

It is widely expected that the next generation of Ethernet will bring a 10X-performance improvement to 10 Gbits/sec. At about the same time the next generation Fibre Channel will only bring a 2x improvement to 2.0 G Bits/sec. This will leave Ethernet 5 times faster than Fibre Channel in the next few years. This is shown in Figure 31.

There is no reason to expect that the high volume production economies of Ethernet will not be continued as shown in the chart above with Ethernet attachment costs being significantly less expensive for both adapters and for switch ports.

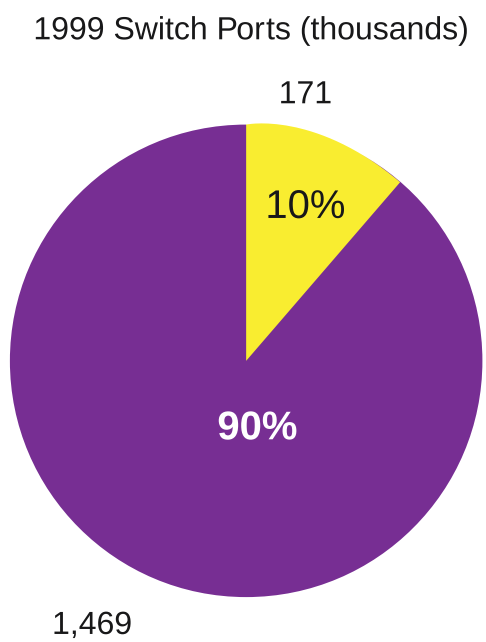

Figure 30 – The continued pervasiveness of Ethernet will continue to result in the lowest possible network infrastructure costs

Target Applications for the NSU

The following table lists applications that are particularly suited for the 3ware NSU in either DAS or E-SAN configurations.

	DAS	E-SAN
Web serving	✔	
Customer Relationship Management	✔	✔
Microsoft Exchange Servers	✔	
Campus OLTP (Online Transaction Processing)	✔	✔
OLAP (Online Analytical Processing)	✔	✔
Nearline - tape replacement	✔	✔
Archival - HSM	✔	✔
"No extra slots disk controller slots"	✔	✔
Compressed audio or video data	✔	✔
Web Cache	✔	
Non Linear Video Editing (NLVE)	✔	
Storage for NAS servers	✔	
Storage for DAS servers	✔	

Table 13 – 3ware target applications and

Technical Specifications of the Network Storage Unit

Supported Operating Systems	Linux 2.2, Win 2K, WinNT

Ethernet Support	10/100/1000 base T
Number of Network Interfaces	2Gb/E, 4x100/E

Power Cords	2
Power Source	Dual load sharing Power Supplies
Rated Ac Input Voltage Requirements	110: 100 - 120 VAC 220: 200 - 240 VAC
Line Frequency	50 - 60 Hz
Current Rating	110: 5A / 220: 2.5A
Number of Power Cords	2
Warranty	1 year

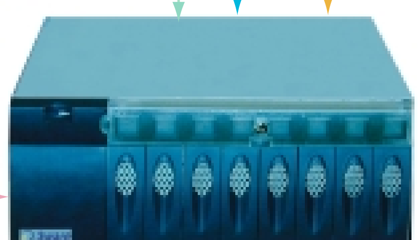

Rack Dimensions HXWXD *(shown)*	6.9"x17.3"x27.4"
Tower Dimensions HXWXD *(not shown)*	19.5'x10"x28.6"
Weight	75 pounds
Operating Temperature	40-95ºF (5-35ºC
Operating Humidity	20-80%
Maximum Wet Bulb	80.6ºF (27.0ºC)
Maximum Operating Altitude	7500 ft

Number of Disk Drive	8
RAID levels supported	0, 1, 10, JBOD
Drive Capacity / Performance	75 GB/7200 RPM
Subsystem Capacity	600 GB Unconfigured

Figure 31 – Technical specifications of the 3ware NSU

Services and Support

3ware offer a one-year warranty on NSU hardware. As a value-added service 3ware offers installation support, certified training, part replacement, annual product upgrades, and support via the Web at www.3ware.com/support, or by phone at the numbers shown on the back of this report. Contact 3ware for details.

A Summary Discussion of Trends and Predictions of the Future

10

The trend to Stack Processing in NIC Hardware

All major Ethernet network interface card (NIC) vendors are implementing techniques to accelerate the processing of the TCP/IP stack on the network interface card to remove this task from the server CPU. This trend is expected to become part of the commodity-pricing trend in this market segment and accelerate the transition to IP-SANs.

FC will double speed in the next two years while Ethernet will increase speed by ten fold in the next two years.

The Fibre Channel consortium originally planned a 2X-speed increase for its next generation, which in retrospect may have been a tactical mistake. This will further enhance the compelling business case of building SANs with readily available Gig/E technology.

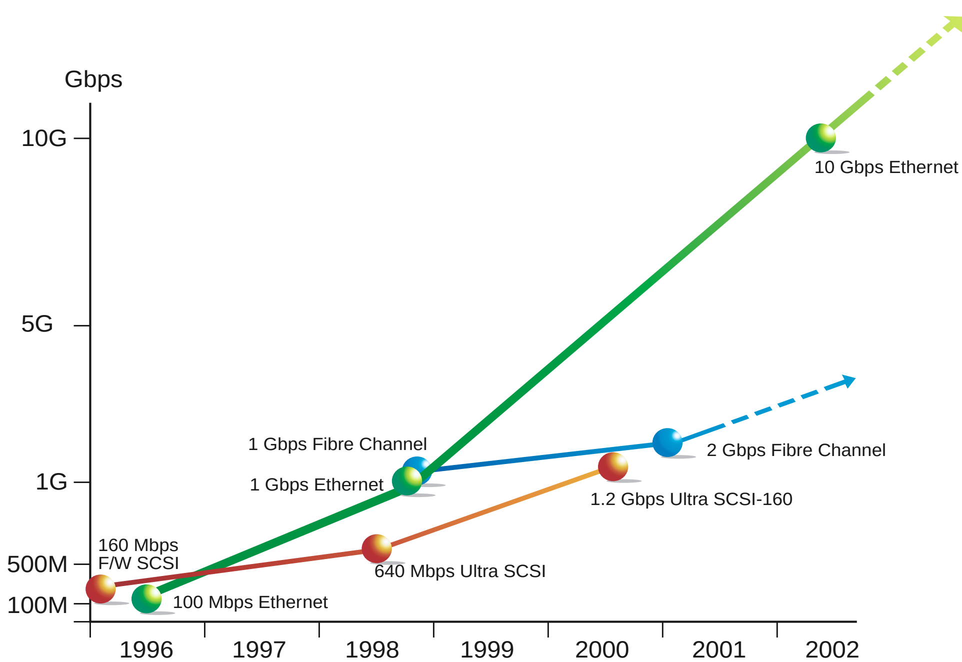

Figure 32 – Understanding bandwidth trends over time and looking into the future

FC-SAN has limited implementations today but many problems

Although there are many problems today with FC-SAN implementations (as shown in Figure 33) it has found a place in large data centers and will most likely coexist with E-SAN in the future.

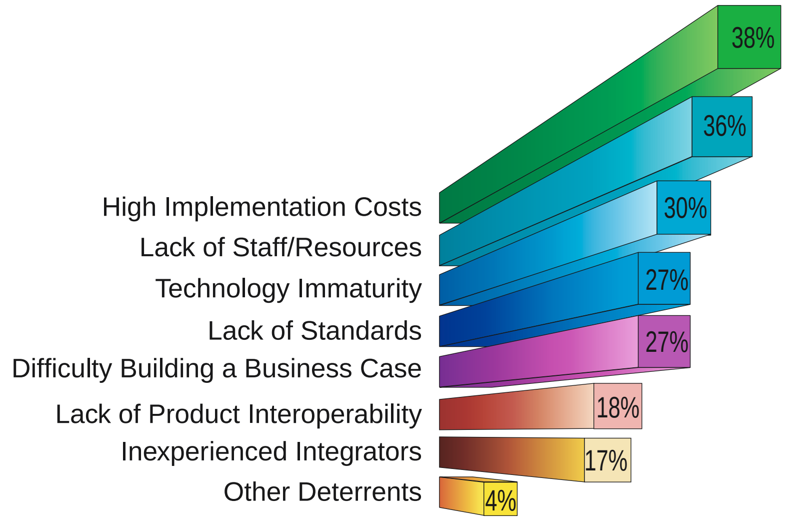

Figure 33 - Deterrents to FC-SAN today

Why FC-SAN deterrents will not apply to E-SAN

Enterprises that do not need to implement SAN today, will wait for the emergence of E-SAN.

Probably the best indicator of the future can be seen Table 13 where the relative advantages and disadvantages is highlighted in blue for each technology. This information provides a clear view that there are major cost, proprietary features, lack of standards, training requirements and capital investment deterrents to Fibre Channel SANs that are not deterrents to E-SAN. These barriers have suppressed the adoption rate of FC-SAN keeping it at about 1 in 10 enterprises. Add to this the 10 kilometer distance limitation of Fibre Channel compared to the effectively unlimited distances of E-SAN by using common carrier facilities which support TCP/IP and the broadly accepted standards for security and congestion control of remote E-SAN. It is very likely that many enterprises that do not need to implement SAN today will wait for the emergence of E-SAN and the many advantages it brings in network speed, security, manageability and cost effectiveness.

Table 14 - A comparison of E-SAN and FC-SAN

	Type of SAN	
Function	**FC-SAN**	**E-SAN**
Current generation bandwidth	1 Gigabit/sec	1 Gigabit/sec
Next generation bandwidth	2 Gigabit/sec	10 Gigabit/sec
Current deployment	8-10% of "lock in"	Emerging technology Deployments expected in 2001
Interoperable switches from multiple vendors	Not today - Proprietary "lock in"	Yes
Distance supported	10 Kilometers	Virtually unlimited distance
Networking protocols	Fibre Channel has a protocol for each layer of the OSI model with (FCP/FC) for the transport and network layers.	TCP/IP
Storage protocol	Fibre Channel	iSCSI or others such as 3ware's Storage Control Protocol (SCP)
Networking fabric	Requires new optical Fibre Channel cabling, switches and network analyzers	Uses existing Ethernet networking infrastructure and standards.
Training investment	Requires new training	Uses networking expertise widely known today.
Storage standards	FC protocol	Looks to the host like standard local SCSI disk.
Integration effort	Integrators are inexperienced and there is much difficulty with most FC-SAN deployments. Vendors have extensive interoperability abs that must test all combinations of components.	Easy to configure and integrate.
Data security	There are Fibre Channel security risks due to lack of standards.	Ethernet provides well understood security protection through well known Network facilities
SNMP Reporting	"Out of band" only	In band or "out of band"
Client access to storage array	Indirect	Direct
Bridge required for distance	Yes	No

Glossary of Terms

3DM

3Ware's Data Management software utility that allows browser based remote management of the NSU over the Web.

10BaseT

Ethernet with a data transfer rate of 10 Mbits/sec or 1.25 Mbytes/sec. Uses Cat3 or Cat5 wiring.

100BaseT

Also known as Fast Ethernet with a data transfer rate of 100 Mbits/sec or 12.5 Mbytes/sec. Uses standard Cat5 wiring.

1000Base-LX/SX

Gigabit Ethernet over fobre optic cable.

1000BaseT

Also known as Gigabit Ethernet over Category 5 copper cable with a data transfer rate of 1000 Mbits/sec or 100 Mbytes/sec.

10,000Base-LX/SX

Also know as 10-Gigabit Ethernet over a fibre optic cable with a data transfer rate of 10,000 Mbits/sec or 1000 Mbytes/sec. Will use fibre optic cabling.

b

Abbreviation for "bit" where 8 "bits" comprise a byte.

B

Abbreviation for byte or the equivalent of one character in text.

Backup Windows

The period of time [usually in the evening] that a computer system is backed up by copying all data to disk or tape.

Backbone

A backbone is another term for bus, the main wire that connects nodes. The term is often used to describe the main network connections comprising the Internet.

Bridge

A bridge is a device that connects two local-area networks (LANs), or two segments of the same LAN. The two LANs being connected can be alike or dissimilar. For example, a bridge can connect an Ethernet with a Token-Ring network. Unlike routers, bridges are protocol -independent. They simply forward packets without analyzing and re-routing messages. Consequently, they're faster than routers, but also less versatile.

Cache

Cache [pronounced cash] can be either a reserved section of main memory or an independent high-speed disk storage device. Two types of caching are commonly used in personal computers: memory caching and disk caching. Disk caching can dramatically improve the performance of applications, because accessing a byte of data in RAM can be thousands of times faster than accessing a byte on a hard disk. When data is found in the cache, it is called a cache hit, and the effectiveness of a cache is judged by its hit rate. The strategies for determining which information should be kept in the cache constitute some of the more interesting problems in computer science. Caching is applied to the Internet whenever a proxy server or web cache is placed at the Internet boundary. The same principles apply as in personal computer caching – that is to store frequently requested data closer to the requestor whether the user is on the internal network or over the web.

Client Server

Client Server Architecture is a network architecture in which each computer or process on the network is either a client or a server. Servers are powerful computers or processes dedicated to managing disk drives (file servers), printers (print servers), or network traffic (network servers). Clients are PCs or workstations on which users run applications. Clients rely on servers for resources, such as files, devices, and even processing power. Another type of network architecture is known as a peer-to-peer architecture because each node has equivalent responsibilities. Both client/server and peer-to-peer architectures are widely used, and each has unique advantages and disadvantages. Client-server architectures are sometimes called two-tier architectures.

CPU

Central Processing Unit. Can refer to either a processor chip such as Sun's SPARC or Intel's Pentium, or to a processor chip or chips and support circuitry on a CPU board. As storage becomes an increasingly larger part of systems design, some have suggested that the definition of CPU be changed to "compute peripheral unit."

Datagram

A datagram is a packet of information carried between a source and a destination using connectionless methods based on routing schemes. IPX (Internetwork Packet Exchange) and Internet Protocol (IP) are datagram services residing in the network layer of the OSI model. A datagram includes the network address, which is examined by routers along the way to determine how to forward the datagram. Datagrams should not be confused with frames. Frames reside in the data link layer of the OSI protocol stack and are addressed to a specific computer on a network segment. A datagram is encapsulated into a frame and is addressed to a specific computer on an internetwork.

Disk Switch

A patented 3ware technology that allows all disk drives in an NSU to avoid the bus arbitration delays of SCSI or the loop arbitration delays of Fibre Channel.

DSS

Decision Support Software.

EBCDIC

Abbreviation of Extended Binary-Coded Decimal Interchange Code. Pronounced eb-sih-dik, EBCDIC is an IBM code for representing characters as numbers. Although it is widely used on large IBM computers, most other computers, including PCs and Macintoshes, use ASCII codes.

E-SAN

Storage Area Networks designed with Ethernet networking protocols of TCP and IP for the transport and networking layers respectively.

ESCON

Enterprise Systems Connection architecture is an IBM mainframe ESA/390 computer peripheral interface or connection between two mainframes for data exchange. The I/O interface utilizes ESA/390 logical protocols over a serial interface that configures attached units to a communication fabric. ESCON is based on networking technology. ESCON provides direct channel-to-channel connections between mainframe systems over fiber-optic links at distances up to 60 kilometers or 36 miles. ESCON also provides a way for communication controllers and other devices to share a single channel to a mainframe.

Ethernet

A Local Area Network (LAN) protocol developed by Xerox in cooperation with Digital Equipment and Intel in 1976. Ethernet supports a star or bus topology and first supported a data transfer rate of 10 megabits per second or 10 Mbps. The Ethernet specification formed the basis of the IEEE 802.3 standard, which specifies the physical and lower software layers. Ethernet uses the CSMA/CD access method for handling simultaneous demands and is one of the most widely implemented LAN standards.

EtherStorage

Adaptec'proposal for SCSI over Ethernet in which they propose a new session layer protocol (SEP) and transport layer protocol (STP) specifically optimized for storage traffic over SANs. See also SEP and STP.

Fast Ethernet

Or 100BaseT, defined by the IEEE 802.3 committee, provides a 100 Mbps standard that is compatible with existing 10BaseT installations, preserving the CSMA/CD media access control (MAC) protocol.

Fibre Channel

Fibre Channel is an ANSI standard designed to provide high-speed data transfers between workstations, servers, desktop computers and peripherals. Fibre Channel makes use of a circuit/packet switched topology capable of providing multiple simultaneous point-to-point connections between devices. Fibre Channel is widely deployed in proprietary SAN implementations today. Standards for Fibre Channel SANs are being worked on by the Storage and Networking Industry Association (SNIA). The technology has gained interest as a channel for the attachment of storage devices, but has limited popularity as high-speed networks interconnect. Fibre Channel can be deployed in point-to-point, arbitrated loop (FC-AL), or switched topologies. Fibre Channel nodes log in with each other and the switch to exchange operating information on attributes and characteristics. This information includes port names and port IDs and is used to establish interoperability parameters.

FC over IP

Fibre Channel over IP is a tunneling protocol for encapsulating FC commands over IP. Also known as FCIP it has been submitted to the IETF. See also SCSI over IP, IETF, IPFC, tunneling.

FCP

An acronym for Fibre Channel Protocol – An ANSI standard covering Fibre Channel protocol for SCSI. See also Fibre Channel Protocol.

FC-SAN

Fibre Channel Sans are Storage Area Networks based on Fibre Channel Protocol (FCP) and FC as the transport and network layer protocols respectively.

FDDI

Fiber Distributed Data Interface. A standard for local area networks that typically uses fiber-optic media capable of data rates up to 100-megabits/second over distances up to 100 km. An FDDI network is a token-based logical ring, and is often constructed as a pair of counter-rotating redundant rings (called dual-attachment mode) for reliability. Ethernet, in contrast, is a bus-based, non-token, 10-megabits/second network standard.

Flow Control

Flow Control mechanisms control the flow of data between sender and receiver so that the receiver's buffers do not overflow.

Frame

In networking, frames are like postal envelopes with addresses called headers. They are sent by routers from one subnetwork to another to their final address. Along the way routers look at the header [address] and forward the frame [letter] onward to its destination much like automated postal scanning equipment scans a zip code and forwards it to the correct post office for delivery.

FTP

FTP is an abbreviation of File Transfer Protocol, the protocol used on the Internet for sending files.

Gigabit Ethernet

Gigabyte Ethernet is a draft standard of the IEEE 802.3 committee which when concluded will provide a mechanism for conveying Ethernet format packets at GB/s speeds. The goals of the gigabit Ethernet effort include: preserve the CSMA/CD access method with support for 1 repeater, use the 802.3 frame format, provide simple forwarding between Ethernet, fast Ethernet and gigabit Ethernet, support both fiber and copper, and accommodate the proposed standard for flow control.

Gigabyte

1024 Megabytes.

GUI

Graphical User Interface refers to the screen design as presented to a user when accessed over the Internet or at a terminal connected to a computer.

HTTP

HTTP is an abbreviation for HyperText Transfer Protocol, the underlying protocol used by the World Wide Web. HTTP defines how messages are formatted and transmitted, and what action Web servers and browsers should take in response to various commands. For example, when you enter a URL in your browser, this actually sends an HTTP command to the Web server directing it to fetch and transmit the requested Web page. The other main standard that controls how the World Wide Web works is HTML, which covers how Web pages are formatted and displayed. HTTP is called a s tateless protocol because each command is executed independently, without any knowledge of the commands that came before it. This is the main reason that it is difficult to implement Web Sites that react intelligently to user input. This shortcoming of HTTP is being addressed in a number of new technologies, including ActiveX, Java, JavaScript and cookies. Currently, most Web browsers and servers support HTTP 1.1. One of the main features of HTTP 1.1 is that it supports persistent connections. This means that once a browser connects to a Web server, it can receive multiple files through the same connection. This should improve performance by as much as 20%.

Hub

A hub is a common connection point for devices in a network. Hubs are commonly used to connect segments of a LAN. A hub contains multiple ports. When a packet arrives at one port, it is copied to the other ports so that all segments of the LAN can see all packets. A passive hub serves simply as a conduit for the data, enabling it to go from one device (or segment) to another. So-called intelligent hubs include additional features that enable an administrator to monitor the traffic passing through the hub and to configure each port in the hub. Intelligent hubs are also called manageable hubs. A third type of hub, called a switching hub, actually reads the destination address of each packet and then forwards the packet to the correct port. Gigabit Ethernet does not support the concept of a hub, and all connections are switched.

InfiniBand

The InfiniBand Architecture is being promoted by seven companies Compaq, Dell, Hewlett-Packard, IBM, Intel, Microsoft and Sun Microsystems. InfiniBand will de-couple the I/O subsystem from memory by utilizing channel-based point to point connections rather than a shared bus, load and store configuration. The newly designed interconnect utilizes a 2.5 Gbit/sec wire speed connection with one, four or twelve wire link widths. This offers scalable performance through multi-link connections as well as a host of interoperable link speeds. The specification will support both copper and fibre implementations.

Internet

Internet refers to a global network connecting millions of computers. As of 1999, the Internet has more than 200 million users worldwide, and that number is growing rapidly. More than 100 countries are linked into exchanges of data, news and opinions.

IETF

The Internet Engineering Task Force (IETF) is a loosely self-organized group of people who make technical and other contributions to the engineering and evolution of the Internet and its technologies. It is the principal body engaged in the development of new Internet standard specifications

Intranet

A network based on TCP/IP protocols (an Internet) belonging to an organization, usually a corporation, accessible only by the organization's members, employees, or others with authorization. An Intranet's Web Sites look and act just like any other Web Sites, but the *firewall* surrounding an Intranet fends off unauthorized access. Like the Internet itself, Intranets are used to share information. Secure Intranets are now the fastest-growing segment of the Internet because they are much less expensive to build and manage than private networks based on proprietary protocols.

IP

The IP (Internet Protocol) is currently in version 4 or IPv4. It is the underlying protocol for routing packets on the Internet and other TCP/IP-based networks. IP is an internetwork protocol that provides a communication standard that works across different types of linked networks for example Ethernet, FDDI or ATM. In an internetwork, the individual networks that are joined are called subnetworks or subnets. IP provides a universal way of packaging information for delivery across heterogeneous subnet boundaries. See also TCP Transmission Control Protocol.

IPFC

This is the term for the IP over Fibre Channel proposal submitted to the IETF for sending network traffic through Fibre Channel switches. It should not be confused with the FCIP proposal for tunneling FC storage traffic over IP. See also FC over IP, SCSI over IP, IETF and tunneling.

IRC

IRC is short for Internet Relay Chat, a chat system developed by Jarkko Oikarinen in Finland in the late 1980s. IRC has become very popular as more people get connected to the Internet because it enables people connected anywhere on the Internet to join in live discussions. Unlike older chat systems, IRC is not limited to just two participants. To join an IRC discussion, you need an IRC client and Internet access. The IRC client is a program that runs on your computer and sends and receives messages to and from an IRC server. The IRC server, in turn, is responsible for making sure that all messages are broadcast to everyone participating in a discussion. There can be many discussions going on at once; each one is assigned a unique channel.

iSCSI

iSCSI is a transport level protocol for the popular SCSI family of protocols to flow over IP as part of the IPS (IP storage) working group of the IETF. The iSCSI protocol is still being defined, but is expected to operate on top of TCP and be fully compatible with the requirements laid out in the SCSI architecture Model - SAM2 document.

L-2, L-3 and L-4 Switches

Switches that operate at level two or three of the OSI reference model are known as L-2 and L-3 switches. An advanced type of switch is called an L4 switch, which operates at the fourth layer (Transport Layer) of the OSI Reference Model and is responsible for the integrity of data transmissions between LAN segments. See also Switch.

LAN

Local area networks or LANs are networks of computers that are geographically close together; this usually means within the same building. Most LANs are confined to a single building or group of buildings. However, one LAN can be connected to other LANs over any distance via telephone lines and radio waves. A system of LANs connected in this way is called a wide-area network (WAN).

Most LANs connect workstations and personal computers. Each node (individual computer) in a LAN has its own CPU with which it executes programs, but it is also able to access data and devices anywhere on the LAN. This means that many users can share expensive devices, such as laser printers, as well as data. Users can also use the LAN to communicate with each other, by sending e-mail or engaging in chat sessions. There are many different types of LANs and Ethernet LANs is the most common for PCs. Most Apple Macintosh networks are based on Apple's AppleTalk network system, which is built into Macintosh computers. LANs are capable of transmitting data at very fast rates, much faster than data can be transmitted over a telephone line; but the distances are limited, and there is also a limit on the number of computers that can be attached to a single LAN.

Load Balancer

Load balancers can prevent servers from being overloaded by distributing large numbers of access requests across multiple Internet servers. The requests are directed to the most-available server based on how busy that server is and other considerations such as storage configuration, server power etc. Load balancers can take the form of special software that runs on a dedicated server, a proprietary black box, or a capability in an advanced LAN switch.

MAC

Media Access Controls or MACs are the rules defined within a specific network type that determines how each station accesses the network cable. Using a token-passing method, a carrier sensing and collision detection method or a demand priority method prevents simultaneous access to the cable. The MAC used for 100BaseT Ethernet as implemented is based on the "demand priority" access method in which the central hub scans all its ports in a round-robin fashion to detect stations that want to transmit a frame. Higher priorities can be requested by ports to transmit real-time information like video or audio.

MAN

Metropolitan Area Networks are networks within a metropolitan area as might for example be used for a city government.

NAS

Network Attached Storage. A LAN-connected storage device that responds to file access requests from a host. A NAS box is actually a small server with its own local file system, that responds to network file requests using either the NFS protocol, typically Unix hosts, or CIFS protocol typically from Microsoft Windows hosts.

NDMP

NDMP is a standard protocol for network-based backup of network-attached storage. NDMP hides the unique interfaces from third party backup software which allows this software to execute on any NDMP compliant system on the network (such as the NS2000 Host Node, and control backups on the NS2000 using standard commands.

NIC

Network Interface Cards (or NICs) in the NS2000 support 10/100BaseT Ethernet, Gigabit Ethernet, FDDI or ATM. There are from one to three on each I/O Node.

Node

See Fibre Channel.

NSU

3ware's Network Storage Unit.

OSI

OSI is an abbreviation for Open System Interconnection, an ISO standard for worldwide communications that defines a networking framework for implementing protocols in seven layers. Control is passed from one layer to the next, starting at the application layer in one station, proceeding to the bottom layer, over the channel to the next station and back up the hierarchy. At one time, most vendors agreed to support OSI in one form or another, but OSI was too loosely defined and proprietary standards were too entrenched. Except for the OSI-compliant X.400 and X.500 e-mail and directory standards, which are widely used, what was once thought to become the universal communications standard now serves as the teaching model for all other protocols. Most of the functionality in the OSI model exists in all communications systems, although two or three OSI layers may be incorporated into one. The model is organized into layers defined as follows:

Level or Layer	Name	Function
7	Application Layer	Application level messages
6	Presentation Layer	Manages data representation conversions. For example, the Presentation Layer would be responsible for converting from EBCDIC (IBM mainframes) to ASCII (UNIX and Windows NT).
5	Session Layer	Responsible for establishing and maintaining communications channels. In practice, this layer is often combined with the Transport Layer.
4	Transport Layer	Responsible for end-to-end integrity of data transmission.
3	Network Layer	Routes data from one node to another.
2	Data Link Layer	Responsible for physical passing data from one node to another.
1	Physical Layer	Manages putting data onto the network media and taking the data off.

Packet

A packet is a piece of a message transmitted over a packet-switching network. One of the key features of a packet is that it contains the destination address in addition to the data. In IP networks, packets are often called ??

PCI

The Peripheral Channel Interconnect is an ANSI standard for an I/O bus used in most PCs and servers.

PDF

PDF stands for Portable Document Format, a file format developed by Adobe Systems. PDF captures formatting information from a variety of desktop publishing applications, making it possible to send formatted documents and have them appear on the recipient's monitor or printer as they were intended. To view a file in PDF format, you need Adobe Acrobat Reader, a free application distributed by Adobe Systems.

Port / Port ID

See Fibre Channel.

RAID

Redundant Array of Independent Disks includes a variety of disk configurations with parity protection against data loss. RAID was first popularized by computer scientists at UC Berkeley in the late 1980's and have become industry standards today.

RDBMS

Relational Database Management Software includes products from Oracle, IBM, Informix and Microsoft.

Reader

A file reader is a software program that reads all the file layout characteristics of a particular format so that it can be displayed on the computer screen.

Router

A Router is a device that connects any number of LANs. Routers use headers and a forwarding table to determine where packets go, and they use ICMP to communicate with each other and configure the best route between any two hosts. Very little filtering of data is done through routers. Routers do not care about the type of data they handle.

SAN

Storage Area Networks [SANs] connect storage subsystems such as disk and tape drives to servers through Fibre Channel or Ethernet switches. SAN storage devices do not have their own file systems, but instead rely on the requesting host's file system to translate the file name and offset into a logical unit number and logical block address. Most SANs today are based on Fibre Channel and are proprietary which is due in large to the lack of a Fibre Channel spec that is being addressed by the SNIA.. EMC's implementation of SAN is known as Enterprise Storage Networks (ESN), Compaq's implementation is called Enterprise Network Storage Architecture (ENSA) and Hitachi's implementation of SAN is known as Freedom Data Networks (FDN). Major benefits of SANs include application transparency, outboard backup, sharing of resources, pooling, and reduced cost of storage management. Storage Area Networks (SAN) are high-speed subnetworks of shared storage devices. SAN architecture works in a way that makes all storage devices available to all servers on a LAN or WAN. As more storage devices are added to a SAN, they too will be accessible from any server in the larger network. In this case, the server merely acts as a pathway between the end user and the stored data. Because stored data does not reside directly on any of a network's servers, server power is utilized for business applications, and network capacity is released to the end user. See also FDN, ENSA, ESN and Fibre Channel.

SCP

3Ware's Storage Control Protocol is a reduced set of commands for the efficient request and retrieval of data from ATA drives in NSU arrays.

SCSI

Small Computer System Interface. An intelligent bus-level interface that defines a standard I/O bus and a set of high-level I/O commands. There are currently many flavors of SCSI defined by different bus widths and clock speeds. The seven major variations of SCSI are SCSI 1, SCSI 2 (Fast / Narrow), SCSI 2 (Fast / Wide), Ultra SCSI (Fast / Narrow), Ultra SCSI (Fast / Wide) – also called SCSI 3, Ultra 2 SCSI (Narrow), Ultra 2 SCSI Wide. Single ended SCSI is used when the peripheral device is close to the point of attachment as in the NS2000 method of attaching disk drives. Differential SCSI provides for reliable operation over greater distances

SCSI over IP

There are a number of proposed and development stage protocols and efforts to flow SCSI block commands over IP. One is the iSCSI spec being developed in the IETF that flows SCSI on top of TCP and IP. Another is the SoIP effort announced by Nishan in June 2000 that flows SCSI on top of UDP and IP. See also FC over IP, SoIP, IETF, and tunneling.

SEP

Adaptec's proposed SCSI over Ethernet session layer protocol that compliments the STP (Storage Transport Protocol) also being developed by Adaptec. See also STP and EtherStorage.

SLIP

SLIP stands for Serial Line Internet Protocol, a method of connecting to the Internet. Another more common method is PPP (Point-to-Point Protocol). SLIP is an older and simpler protocol, but from a practical perspective, there's not much difference between connecting to the Internet via SLIP or PPP. In general, service providers offer only one protocol although some support both protocols. See also PPP.

SNIA

The Storage Networking Industry Association (SNIA) is a not-for-profit organization, made up of 150 companies and individuals spanning virtually the entire storage industry. SNIA members share a common goal, to set the pace of the industry by ensuring that storage networks become efficient, complete, and trusted solutions across the IT community. To this end the SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. The SNIA is working on standards for Fibre Channel SANs and for Network Data Management Protocol. See also NDMP, Fibre Channel.

SoIP

SCSI over IP or SoIP is an effort announced by Nishan in June 2000 that flows SCSI on top of UDP and IP. See also SCSI over IP, iSCSI, FC over IP, IETF and tunneling.

Stack

1) In programming, a special type of data structure in which items are removed in the reverse order from that in which they are added, so the most recently added item is the first one removed. This is also called last-in, first-out (LIFO). Adding an item to a stack is called pushing. Removing an item from a stack is called popping. (2) In networking, short for protocol stack.

STP

SAN Transport Protocol (STP) is a lightweight transport protocol being developed by Adaptec for EtherStorage [Adaptec's term for SCSI over IP] that addresses the poor performance of the TCP/IP stack and processor utilization problems of E SAN by optimizing the transport layer for SAN environments and storage traffic. STP is proposed along with a new session layer protocol called SEP (SCSI over Ethernet Protocol). See also EtherStorage and SEP.

Switch

In networking terminology, a switch is a computing device that filters and forwards packets between Local Area Network (LAN) segments. Switches operate at the data link layer (layer 2) of the OSI Reference Model and therefore support any packet protocol. A special type of switch call an L4 switch operates at the fourth layer (Transport Layer) of the OSI Reference Model and is responsible for the integrity of data transmissions between LAN segments. LANs that use switches to join segments are called switched LANs or, in the case of Ethernet networks, switched Ethernet LANs.

T1

T1 refers to a dedicated phone connection supporting data rates of 1.544Mbits per second. A T-1 line actually consists of 24 individual channels, each of which supports 64Kbits per second. Each 64Kbit/second channel can be configured to carry voice or data traffic. Most telephone companies allow you to buy just some of these individual channels, known as fractional T-1 access. T-1 lines are a popular leased line option for businesses connecting to the Internet and for Internet Service Providers (ISPs) connecting to the Internet backbone. The Internet backbone itself consists of faster T-3 connections. T-1 lines are sometimes referred to as DS1 lines.

T3

T3 A dedicated phone connection supporting data rates of about 43 Mbps. A T-3 line actually consists of 672 individual channels, each of which mainly Internet Service Providers (ISPs) use supports 64 KBPS. T-3 lines for connecting to the Internet backbone and for the backbone itself. T-3 lines are sometimes referred to as DS3 lines.

TB

A Terabyte (TB) equals 1024 Gigabytes.

TCP

Transmission Control Protocol or TCP is a transport layer component of the Internet's TCP/IP protocol suite. It sits above IP in the protocol stack and provides reliable data delivery services over connection-oriented links. TCP uses IP to deliver information across a network and makes up for the deficiency of IP providing a guarantee of reliable delivery services that IP does not. TCP messages and data are encapsulated into IP datagrams and IP delivers them across the network.

TIFF

TIFF is an acronym for tagged image file format, one of the most widely supported file formats for storing bit-mapped images on personal computers (both PCs and Macintosh computers). Other popular formats are BMP and PCX. TIFF graphics can be any resolution, and they can be black and white, gray-scaled, or color. Files in TIF format often end with a .tif extension.

Tunneling

A tunnel is used to deliver data packets across networks that use different protocols than the source and destination network. For example the IETP Fibre Channel over IP workgroup is defining a method of encapsulating Fibre Channel packet into an IP packet and then delivered across the long haul IP networks.

TwinStor

Patented 3ware technology that provides for high-performance mirroring of disks in the NSU arrays.

WAN

Wide Area Networks or WANs are networks of computers that are geographically dispersed and connected by radio waves, telephone lines or satellites.

WAV

The defacto standard for storing sound in files that was developed by IBM and Microsoft.

Web

A system of Internet servers that support specially formatted documents. The documents are formatted in a language called HTML (HyperText Markup Language) that supports links to other documents, as well as graphics, audio, and video files. This means you can jump from one document to another simply by clicking on hot spots. Not all Internet servers are part of the World Wide Web. There are several applications called Web browsers that make it easy to access the World Wide Web. Two of the most popular being Netscape Navigator and Microsoft's Internet Explorer.

Web Site

A site (location) on the World Wide Web. Each Web site contains a home page, which is the first document users see when they enter the site. The site might also contain additional documents and files. Each site is owned and managed by an individual, company or organization.

Webmaster

An individual designated at a Web Site to be in control of many of the operational aspects of the site.

3ware, Inc.
701 E. Middlefield Road
Suite 300
Mountain View, CA 94043
T 650.327.8600
F 650.254-5690
sales@3ware.com
www.3ware.com

East Coast Office
T 919.460.9282
F 919.460.8968
sales@3ware.com

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by 3ware, Inc. This document describes some capabilities that may be configuration-dependent, and features that may not be currently available.

The contents of this document represents the opinion of 3ware Inc. and includes accurate information as far as can be determined at the time of publication. 3ware, Inc. is not responsible for inaccuracies resulting from changing technologies

Contact the nearest 3ware, Inc sales office for information on feature and product availability.