

For more PPT's and PDF's Log on to Studymafia.org

A
Seminar
Report
On
Storage Area Network (SAN)

Submitted in partial fulfillment of the requirement for the award of degree
Of Bachelor of Technology in Computer Science

SUBMITTED TO:

www.studymafia.org

SUBMITTED BY:

www.studymafia.org

ACKNOWLEDGEMENT

It is my pleasure to be indebted to various people, who directly or indirectly helped me in the seminar on “Storage Area Network”.

I would like to express my gratefulness to....., who has given me the opportunity to carry out this seminar.

Lastly, I would like to thank the almighty and my parents for their moral support.

INTRODUCTION

The recent explosion in e-business activity and internet commerce has provided organization with unlimited opportunities for developing new information delivery channel. Data is today perceived to be the key asset for many organizations such as banking, stock exchange, government record etc. This has generated an explosive demand for data storage and this demand can be addressed by deploying SAN. The activity to share a single large storage device across many server or application has made SAN an attractive option in today's market place.

As organization continue to broaden there reach to business partners and customers around the globe, they expose key IT system to a wider range of potential security threats. Today data theft, fraud, hacker attempts, and human error increasingly threaten security of information exchange within the enterprise and across the public networks, such as the internet. In order to protect the key data stored, storage networking venders are rapidly deploying and developing security frameworks that help ensure safe reliable data processing throughout a storage area networks(SAN).

The most common thing to remember about SAN security is that SAN is a network and is vulnerable to the same sorts of vulnerabilities and attacks that more conventional computers are. SAN resources can be protected by physical security and the hosts on the SAN should be expected to meet stringent security requirements. As SAN continue to grow, it will become a bigger target for malicious attackers. In this section we will examine the emergence and evolution of the fiber channel protocol in the context of storage network technology.

WHAT IS A STORAGE AREA NETWORK

According to the Storage Networking Industry Association dictionary (and who should know better?):

A storage area network (SAN) is any high-performance network whose primary purpose is to enable storage devices to communicate with computer systems and with each other.

We think that the most interesting things about this definition are what it *doesn't* say:

- *It doesn't say that a SAN's only purpose is communication between computers and storage.* Many organizations operate perfectly viable SANs that carry occasional administrative and other application traffic.
- *It doesn't say that a SAN uses Fiber Channel or Ethernet or any other specific interconnect technology.* A growing number of network technologies have architectural and physical properties that make them suitable for use in SANs.
- *It doesn't say what kind of storage devices is interconnected.* Disk and tape drives, RAID subsystems, robotic libraries, and file servers are all being used productively in SAN environments today. One of the exciting aspects of SAN technology is that it is encouraging the development of new kinds of storage devices that provide new benefits to users. Some of these will undoubtedly fail in the market, but those that succeed will make lasting improvements in the way digital information is stored and processed.

Let's dig a little deeper into this definition.

What Makes a SAN Different?

Anyone in the information technology field knows very well that computers are already connected to storage devices. If that's all a SAN does, what's new or different about it? The answer is a simple phrase that we'll keep coming back to over and over throughout this book:

Universal Storage Connectivity

Computers are indeed connected to storage today, but are all of an installation's computers connected to all of its storage? That's the key point about SANs—they connect lots of computers to lots of storage devices, enabling the computers to negotiate device ownership among themselves and, ideally, to share data. If there is one defining characteristic of a SAN, its universal connectivity of storage devices and computers.

To appreciate the value of universal storage connectivity, consider the conventional client/server computer system depicted in Figure 1.1.




Figure 1.1 Client/Server information islands

From this business-as-usual picture of client/server computing, it's immediately apparent that by deploying multiple servers, an organization automatically creates unconnected islands of information. Each island is accessible by one computer but not the others. If Computer B needs to use data that was produced by Computer A, that data has to be copied to Computer B.

There are several techniques for moving data between computers: **backup, file transfer; and interprocess communication**, to name a few. But the real issue is that the information services organization has to acquire and manage the extra resources required both to copy data from Computer A to Computer B and to store it at both sites. There's no business reason for this duplication of effort, other than that a computer needs data that was produced by another computer.

There's a more serious implication of an information processing strategy that relies on regular copying of data from computer to computer. Computers that receive data copies are often forced to work with data that is out of date simply because it's physically impossible to make copies in a timely fashion. Moreover, the extra operational complexity introduced by having to copy data between servers creates additional opportunity for costly errors.

Contrast this with the SAN based distributed system architecture illustrated in Figure 1.2.




Figure 1.2 A SAN eliminates islands of information

With a SAN, the concept of a single host computer that owns data or storage isn't meaningful. All of the servers in Figure 1.2 are physically connected to all of the storage devices. If Server F needs data that was produced by Server E, there's no need to copy it, because Server F can access the devices on which Server E stored the data. All that's required is a logical change of storage device ownership from Server E to Server F or better yet, an agreement by Server E to stay out of the way while Server F is actively using the data.

Universal storage connectivity has some pretty powerful implications for information services departments:

- There's no need to devise mid schedule data transfers between pairs of servers.
- There's no need to purchase and maintain extra storage for temporarily staging one server's data at another server.
- There's no need to worry about whether copies of data being used by two computers running different applications are synchronized (i.e., have exactly the same contents), because the two computers are working from the same copy of data.

Indeed, at this simplistic level, it's hard to see how any organization responsible for electronic information processing could not want a SAN. Let's drill still deeper and see how true that is.

What Makes a Good SAN?

The completely interconnected SAN architecture shown on the right side of Figure 1.2 is intuitively attractive, but if it's going to be the I/O backbone of an information services operation, it needs to have couple of qualities:

A SAN must be highly available. A single SAN connecting all computers to all storage puts a lot of enterprise information accessibility eggs into one basket. The SAN had better be pretty indestructible or the enterprise could literally be out of business. A good SAN implementation will have built-in protection against just about any kind of failure imaginable. As we will see in later chapters, this means that not only must the links and switches composing the SAN infrastructure be able to survive component failures, but the storage devices, their interfaces to the SAN, and the computers themselves must all have built-in strategies for surviving and recovering from failures as well.

The I/O performance of a SAN must grow or scale as the number of interconnected devices grows. If a SAN interconnects a lot of computers and a lot of storage, it had better be able to deliver the performance they all need to do their respective jobs simultaneously. A good SAN delivers both high data transfer rates and low I/O request latency. Moreover, the SAN's performance must be able to grow as the organization's information storage and processing needs grow. As with other enterprise networks, it just isn't practical to replace a SAN very often.

On the positive side, a SAN that does scale provides an extra application performance boost by separating high-volume I/O traffic from client/server message traffic, giving each a path that is optimal for its characteristics and eliminating cross talk between them.

The investment required to implement a SAN is high, both in terms of direct capital cost and in terms of the time and energy required to learn the technology and to design, deploy, tune, and manage the SAN. Any well-managed enterprise will do a cost-benefit analysis before deciding to implement storage networking. The results of such an analysis will almost certainly indicate that the biggest payback comes from using a SAN to connect the enterprise's most important data to the computers that run its most critical applications.

But its most critical data is the data an enterprise can least afford to be without. Together, the natural desire for maximum return on investment and the criticality of operational data lead to Rule 1 of storage networking:

When designing a SAN to access critical enterprise data, make sure the SAN is highly available (i.e., can survive failures of both components in it and components attached to it) and make sure it can grow well beyond anticipated peak performance needs without disruption.

What Makes a Great SAN?

So the fundamental feature of SANs is universal data storage connectivity. Universal connectivity enables a host of important benefits. Depending on the particular SAN hardware and software components chosen, additional benefits may accrue from advanced functions being built into today's SAN devices. Again, we describe some specific features and benefits later on, but for now we assert Rule 2 of storage networking:

When evaluating SAN implementation options, once the basic capacity, availability, and performance requirements can be met, look for advanced functionality available in the chosen architecture and consider how it might be used to further reduce cost or enhance the information services delivered to users.

WHY CONNECT STORAGE TO A NETWORK?

Throughout the journey into storage networking, it's important to keep sight of the benefits being sought. Throughout this book we try hard to distinguish between the features of storage networking, such as universal connectivity, high availability, high performance, and advanced function, and the benefits of storage networking that support larger organizational goals, such as reduced cost and improved quality of service.

The specific benefits that storage networking delivers are different in every situation, but with storage networking, as with any other aspect of information technology, benefits can be broadly classified as either:

- Reducing the cost of providing today's information services or providing or enabling new services that contribute positively to overall enterprise goals.

Storage networking offers ample opportunity for an information services department to deliver both types of benefits. For example, in the realm of cost savings:

- If all online storage is accessible by all computers, then no extra temporary storage is required to stage data that is produced by one computer and used by others. This can represent a substantial capital cost saving.
- Similarly, if tape drives and robotic media handlers can be accessed directly by all computers, fewer of these expensive and infrequently used devices are needed throughout the enterprise. This, too, reduces total enterprise capital cost for information processing without diminishing the quality of service delivered.
- Probably most important, however, are the administrative and operational savings in not having to implement and manage procedures for copying data from place to place. This can greatly reduce the cost of people—the one component cost of providing information services that doesn't go down every year!

Similarly, consolidating storage on a network may enable information services departments to provide services that just aren't possible with storage attached directly to each computer. For example:

- SAN connectivity enables the grouping of computers into cooperative clusters that can recover quickly from equipment or application failures and allow data processing to continue 24 hours a day, every day of the year.
- With long-distance storage networking, 24×7 access to important data can be extended across metropolitan areas and indeed, with some implementations, around the world. Not only does this help protect access to information against disasters; it can also keep primary data close to where it's used on a round-the-clock basis.
- SANs remove high-intensity I/O traffic from the LAN used to service clients. This can sharply reduce the occurrence of unpredictable, long application response times, enabling new applications to be implemented or allowing existing distributed applications to evolve in ways that would not be possible if the LAN were also carting I/O traffic.
- A dedicated backup server on a SAN can make more frequent backups possible because it reduces the impact of backup on application servers to almost nothing. More frequent backups means more up-to-date restores that require less time to execute.

Figure 1.3 illustrates the role a SAN occupies in a distributed client/server computing network.




Figure 1.3 A SAN separates client traffic from disk I/O

THE SECRET TO SANS' SUCCESS: SOFTWARE

Today, most of the public attention given to SANs is focused on the interconnects (such as Fibre Channel) that allow universal storage connectivity and the storage devices and computers that connect to them. But interconnects by themselves don't add any functionality to information processing, they only enable functionality to be added. To realize the benefits promised by SANs, not only must the hardware connectivity, performance, availability, and function be in place to enable them, system and application software must take advantage of the hardware to deliver them. Thus we have Rule 3 of storage networking:

Hardware makes SANs possible; software makes SANs happen.

SAN Software Capability

When evaluating SAN technology, the hardware components deserve close scrutiny, to be sure. More important, however, one must also scrutinize the software capabilities carefully to ensure that the implementation will deliver the functionality enabled by the hardware. The following sections give some examples of how software can make a SAN sing.

Sharing Tape Drives

It's very reasonable to expect to share a SAN-attached tape drive among several servers because tape drives are expensive and they're only actually in use while back-ups are occurring. If a tape drive is connected to computers through a SAN, different computers could use it at different times. All the computers get backed up. The tape drive investment is used efficiently, and capital expenditure stays low.

There's just one tiny problem. What is to keep a computer from (accidentally) writing to a tape while another computer is doing a backup? For two or three computers, an administrator can personally schedule tape drive usage so that this doesn't happen (unless something malfunctions, of course). As the number of computers and tape drives grows or as the applications running on them change, manual scheduling becomes increasingly difficult. What's really needed is a foolproof way for computers to negotiate for exclusive ownership of tape drives for the duration of a backup. When one computer wins a negotiation and is using a tape drive, others must be fenced off and kept from using the drive, even if an errant application or faulty backup schedule tries to do so.

Shoring Online Storage Devices

Sharing online storage, as shown in Figure 1.4, that's housed in an enterprise RAID subsystem is similar to sharing tape drives, except that more of it goes on mid requirements for configuration changes are more dynamic. A typical enterprise RAID subsystem makes the online storage capacity of one or more arrays of disks appear to be one or more very large, very fast, or very reliable disks. For now, accept that a RAID subsystem can look like several virtual disks from the viewpoint of host servers. It is quite reasonable that different servers be able to access those virtual disks at different times. For example, one server might collect a business day's transaction records on disk and hand them off to another server at the end of the day for summarization, analysis, or backup.

The problem here is similar to the problem of sharing SAN-attached tape drives among several servers: A server cannot be permitted to write data to a disk that another server owns at the moment. Some foolproof way of fencing off, or preventing other servers from accessing an allocated virtual disk, is needed. Moreover, a distributed system needs to remember which virtual disks have been allocated to (are owned by) which servers, even when one or more of the servers is shut down and rebooted.




Figure 1.4 Coordinating shared devices in a SAN

Again, it's software to the rescue. Some SAN infrastructure implementations help by subdividing a SAN into disjoint zones, and some RAID subsystems allow virtual disks to be reserved as the private property of a single computer (masked from other servers). It takes management software to exploit all these features, or to simulate them in SAN implementations whose components don't offer the features.

Application Failover




Figure 1.5 Application failover in a SAN-based cluster

Although the concept sounds simple, the clustering of two or more computers to enhance application availability is actually one of the most complex operating system problems there is. Even recognizing and isolating failures is nontrivial. Once an application failure is recognized, ownership of exactly the resources it requires to run (and no others) has to be transferred to a failover server immediately, completely, and unambiguously. The application must be restarted and its state at the time of the failure recreated as nearly as possible. Again, it's software in the form of a cluster server or cluster manager that does all this.

Sharing Data

More advanced forms of computer clustering, as illustrated in Figure 1.6, allow for the concurrent sharing of data among different applications running on different servers. This can be extraordinarily useful, for example, for incremental application growth or scaling. Simply stated, if an application outgrows the server on which it is running, don't replace the server with a bigger one. Instead, simply connect another server with the necessary incremental power to the SAN, leaving the original system in place. Both servers can run separate copies, or instances, of the application, processing the same copy of data. More servers can be added as application capacity requirements grow.




Figure 1.6 Shared data clusters enable applications to scale.

Applications configured this way are inherently highly available. If one of several servers running instances of the same application processing the same data falls, the remaining servers just take up the slack. Overall performance may decrease, but at least the application keeps running. That's much better than the alternative.

Data sharing can take any of several forms, from concurrent read-only access to different files in the same file system, all the way to the sharing of a single database by multiple instances of a single application. In all cases, however, the mechanisms for orderly sharing—keeping different applications from overwriting each others' updates to data—are provided by... you guessed it... software.

Direct Data Movement between Devices

Earlier, we mentioned the emergence of more intelligent storage devices with new functionality that is accompanying SAN deployment. One of the key functions being built into intelligent storage devices is the ability to exchange data directly with other intelligent devices in response to commands from a third party (e.g., a computer). Figure 1.7 illustrates such a set up. This capability promises to revolutionize backup, for example, by halving the amount of data movement required to create any kind of backup copy. Less obvious, but potentially at least equally useful, are the possibilities of mirroring or replicating data directly from one device to others.

Direct movement of data between devices needs two kinds of software. First, it needs software or firmware that implements the capability within the devices themselves. Second, host computer applications like backup and replication management that use the capability are required.




Figure 1.7 Direct copy of data between SAN devices

Summarizing SANs and Software

So, while high-performance, highly available, fully interconnected SANs are necessary to realize the benefits we have discussed, they are hardly sufficient. It takes system software on a massive scale to bring to life the benefits that are promised by the underlying properties of SANs. The software required to make the most of SANs takes two forms:

- *System applications.* These applications build upon the basic SAN properties to provide a functionally enhanced execution environment for business applications. System applications include clustering, data replication, direct data copy between devices and the utility functions that use it, and so forth.
- *Management applications.* These applications manage the inherently more complex distributed system environment created by the presence of SANs. Zoning, device discovery, allocation, and RAID subsystem configuration are examples of applications that fall into this category.

Both of these types of application are necessary for the success of SANs and, ultimately, for their utility as a tool for reducing cost and improving the quality of enterprise information services. While this book discusses the hardware components that constitute SANs, a good portion of it is also given over to a discussion of the software capabilities that make SANs into such extraordinarily useful business data processing tools.

THE BEST IS YET TO COME: RADICAL CHANGES IN INFORMATION STORAGE AND PROCESSING

Storage area networks are being deployed in large numbers today and are being used to solve real business problems. But, compared to their ultimate potential, the surface is just being scratched. This section touches on some of the more radical changes in distributed information storage and processing enabled by SANs.

Best-in-Class Computer Systems

SANs enable the interconnection of many servers to many storage devices of different types, including disk arrays and robotically controlled tape drives. Who is it that makes these devices and how is it that they all work correctly together?

Both of the market-leading SAN interconnect technologies, Fibre Channel and Ethernet, as well as the emerging Infiniband, are standards-based—that is, their specifications are defined by voluntary associations of (sometimes competing) companies and are easily available to the general public, rather than being the private property of one company.

Since the same interconnect specifications are available to anyone, in principle, any company should be able to build devices that interoperate with any other company's complementary devices, creating heterogeneous storage area networks. This, in fact, is a stated goal of the storage networking vendor community.

This goal has only partly been realized today. No vendor has developed all the components required to build a complete SAN, but most vendors are engaged in partnerships to qualify and offer complete SANs consisting of the partners' products. Today, consumers are well advised to make SAN purchases from a single vendor or, at a minimum, to choose a primary vendor for major components such as RAID subsystems and accept that vendor's recommendations for other SAN components. The industry is working toward a mode of operation similar to that of the LAN industry, where component interactions are sufficiently well understood that users feel relatively comfortable making purchases at the component level. Much energy is being devoted to interoperability, both in terms of making standards more precise and in terms of validating components against standard test suites rather than against individual complementary implementations.

Smarter Storage and Appliances

While the cost and performance of online storage have improved pretty dramatically over the last decade, its functionality has remained pretty much the same. While other major computer system subassemblies, such as network interface cards and video controllers, have become more intelligent and increasingly autonomous, a disk has pretty much remained a disk and a tape has remained a tape for the last decade.

It's only partly coincidental that with SANs becoming popular, storage device functionality is changing as well. Today, intelligent storage devices that can transfer data among themselves without passing it through a setter are starting to be delivered. Peer-to-peer I/O improves quality of service by halving the bandwidth required to do bulk data transfers such as backup or replication of large file systems and databases. This will ultimately enable application designs that rely more on bulk data transfers, for example, to create more frequent backups or data replicas.

With intelligent storage devices, other advanced functions are also possible. For example, server-based hierarchical storage management software that compresses infrequently used data is available today. Implementing this capability in an enterprise RAID subsystem is technically feasible.

For other examples of how storage device intelligence might be exploited in the future, one need look no further than the storage middleware segment. Backup within the storage subsystem, clusters of storage subsystems, and autonomous global data replication are all possible. Some vendors have introduced these capabilities in storage appliances that are nothing more than computers dedicated to managing large amounts of storage and I/O capacity. With these storage appliances, the distinction between a computer and an enterprise storage subsystem essentially comes down to whether or not the box runs applications.

Heterogeneous Computer Systems

Today, it is quite common for users to partition the storage in large RAID subsystems mid allocate the partitions to different computers made by different vendors and running different operating systems. The computers share the RAID subsystems' common resources (power, cooling, processors, internal bandwidth, cache, etc.), but each has its own private online storage capacity.

With a SAN, however, it should also be possible to share storage and data across computers. In principle, computers from Sun, Hewlett-Packard, IBM, and others should all be able to access data stored in one big enterprise storage subsystem made by EMC, Hitachi Data Systems, HP, or any other enterprise storage subsystem vendor.

Today, data sharing among heterogeneous computers is possible with network attached storage or NAS devices. With other network storage devices, most notably enterprise RAID subsystems, it is not generally possible to share data among computers of different types (although some clustering technologies support data sharing among computers of the same type).

Many difficult technical problems must be solved before different types of computers can share data. By far the greatest of these is that each operating system and file system uses its own unique disk format, simply because there has never been any motivation to do otherwise. SANs provide this motivation because they interconnect the storage of different types of computers running different operating systems. Today, therefore, both academic and industrial researchers and developers are busily working toward the goal of universally accessible online data files.

Data Storage as a Service

Online information storage has traditionally been something that enterprises have implemented and managed for themselves. Whether so inclined or not, information services professionals have been forced to develop storage expertise, as it is they who are responsible for safeguarding the organization's information assets and keeping them accessible.

SANs are enabling a new generation of storage service providers, companies whose business consists of storing and managing the data from other enterprises. These companies take on the hard problems of providing secure data access at contractually defined quality-of-service levels. In the not-too-distant future, many enterprises may literally get their data storage from a plug in the wall, much as they get their voice and data communications services from external providers today.

Similarly, the long-distance connectivity of SANs will someday enable the secure inter-connection of storage networks belonging to enterprises that do business together or that share a specific market. Storage intranets will allow partnering enterprises to replicate or even share information without risk of its being divulged to (possibly hostile) outside parties.

Widespread adoption of these new modes of operation would clearly require significant technological development security, as well as changes in user attitudes. They may or may not

become prevalent. The point to be made here, however, is that SANs are enabling developers, entrepreneurs, and users to break out of the box as they design enterprise information storage and processing infrastructures. There are all kinds of possibilities. We probably haven't even envisioned the ones that will be the success stories of five years from now.

BACK TO EARTH ...

But the information technology professionals in most organizations are responsible for keeping their enterprises' data secure and accessible while they await the brave new world of information processing. Even from this short-term vantage point, SANs have tremendous potential. Used optimally, SANs can improve applications' access to data, dramatically reduce unnecessary redundancy, eliminate bandwidth contention between client messages and data access, and off-load resource-intensive tasks like backup to intelligent storage devices on I/O-optimized interconnects. Information technology professionals cannot afford to ignore the near-term benefits of SAN technology.

A COUPLE OF CLOSING CLARIFICATIONS

This section touches upon a couple of aspects of terminology that in the authors' experience tend to be sources of confusion rather than clarification in discussions of storage networking. We hope it dispels some of the very basic uncertainty around SAN terminology for readers and lets them get on to the serious stuff.

Is a SAN Storage or a Network?

Strictly speaking, a SAN is a storage area network. Much energy is expended uselessly on silly arguments about whether the term SAN should include storage and computers or only the interconnects and other components that link the storage and computers together. Since it's difficult to imagine talking about a SAN without talking about the storage and computers it interconnects, we adopt the more inclusive usage in this book:

A storage area network (SAN) is a network of interconnected computers and data storage devices. As we use it in this book, the term SAN specifically includes both the interconnection infrastructure and the computers and storage that it links together.

Of course, we frequently have occasion to refer specifically to the storage devices attached to a SAN, as distinguished from the more inclusive concept. We use the terms SAN-attached storage and SAN-attached storage device for this purpose.

What Is This Area Thing, Anyway?

One might wonder what the area in a storage area network is, and rightly so. From what we've described so far, the term storage network would seem to apply equally well if not better. What is this area thing, anyway?

There's a strong relationship between the SAN concept and the popular I/O interconnect technology called Fibre Channel. In fact, it's fair to say that Fibre Channel technology is what got the whole SAN movement started.

In the mid-1990s, the companies committed to developing Fibre Channel were in the doldrums. Five years or more had been spent on development, and the products and market weren't emerging as fast as investors would have preferred. Something was needed to generate excitement among consumers and jump-start growth.

Members of the Fibre Channel Loop Community and the Fibre Channel Association, the two trade groups promoting the technology at that time, hit upon the idea of capitalizing on the similarities between Fibre Channel and local area networking:

Universal connectivity. With Fibre Channel, lots of storage devices can be connected directly to lots of computers.

Campuswide separation. With Fibre Channel, storage devices don't have to be in the computer room. They can be anywhere on a fairly sizable campus.

Bandwidth multiplication. With switches (Chapter 7 describes switch capabilities), many high-performance links can be aggregated into a fabric with even higher aggregate performance.

Dynamic reconfiguration. Fibre Channel allows devices and computers to be connected to and disconnected from the network while it is operating.

Enabling of new storage usage styles. Just as local area networks enabled client/server computing, Fibre Channel connectivity enables new ways of using storage to achieve business goals. Sharing access to disk and tape subsystems and direct copying of data between devices are two examples.

Noticing the networking industry's success in creating a strong identity for itself through the use of the LAN acronym, the Fibre Channel community set out to achieve something similar by introducing the term storage area network, or SAN, to describe networks of intelligent storage devices and servers. White papers were written, lectures were given, and articles were published, all pointing out similarities between the local area network that connects all the computers in a building or campus and a storage area network, with the potential to likewise connect all the storage within a similar area. The SAN acronym, with its implicit analogy to LANs, was repeated over and over again.

Although the SAN concept was largely a so-called "marketecture" when first introduced, the aforementioned promises have essentially been delivered upon in the intervening years. The perception created by early Fibre Channel industry marketers has become reality. Products and capabilities are being delivered and the term SAN has become thoroughly embedded in the public consciousness.

Besides, storage networking makes a lousy acronym.

Is "NAS" Just "SAN" Spelled Backward?

One often hears discussions of network-attached storage, or NAS, usually in contexts that leave unclear whether NAS is a kind of SAN, a competitor to SAN technology, or something else entirely.

The term *network-attached* storage (NAS) is used pretty consistently throughout the industry to denote intelligent storage devices that connect to networks and provide file access to clients, which may range from desktop computers to enterprise-class application and database servers. From the time of their introduction, NAS devices have almost universally used Ethernet interconnects and TCP/IP-based protocols to connect to and communicate with their clients.

As we mentioned earlier, the term SAN is typically understood as including the storage attached to the network. What we didn't say earlier is that today the term SAN is also closely identified with block-access storage devices that is, disk drives and devices such as RAID subsystems that behave as disk drives from the client's viewpoint.

Thus, as the terms are commonly used today:

A NAS device provides file access to clients to which it connects using file access protocols (primarily CIFS and NFS) transported on Ethernet and TCP/IP.

A SAN device (or, as we promised to call it, a SAN-attached storage device) is a block-access (i.e., it is a disk or it emulates one or more disks) that connects to its clients using Fibre Channel and a block data access protocol such as SCSI.

Figure 1.8 illustrates these two storage paradigms.




Figure 1.8 SAN and NAS storage devices

A glance at Figure 1.8 clarifies the major architectural difference between NAS and SAN storage as the terms are commonly used today. In the NAS paradigm, the file system that organizes blocks of storage into objects that are convenient for applications to deal with resides in the storage device. The NAS storage device is responsible for allocating storage space and for keeping clients from stepping on each others' toes as they make (possibly conflicting) file access requests. On the host side of the interconnect, a file access client translates applications' file I/O requests into network messages and sends them to the NAS device for execution.

By contrast, in today's SAN paradigm, the file system is on the computer side of the interconnect. Systemwide storage capacity management and conflicts among client data access requests are resolved by cooperation among the SAN-attached servers. This makes host-side software much more complex than with NAS devices.

By absorbing the file system (and hence the on-disk data format) into the storage device, the NAS model makes concurrent data access by different types of computers easy. In fact, today, NAS is the only widely available way to make the same data accessible to computers of different types.

That doesn't quite tell the whole story. NAS file access protocols are very general and functionally rich. Moreover, they usually connect to TCP/IP-based networks, which are designed to support very general interconnection topologies. Because of their functional richness and generality, these protocols are predominantly implemented in software, which executes slowly compared to the device-level firmware and hardware typically used to implement SAN protocols. Raw data access performance of NAS devices, therefore, tends to be lower than that of

otherwise comparable SAN devices, and both client and server processor utilization for accessing data tends to be higher. In simple terms, the trade-off today is, therefore, as follows:

- *Choose NAS* for simplicity of data sharing, particularly among computers and operating systems of different types.
- *Choose SAN* for the highest raw I/O performance between data client and data server. Be prepared to do some additional design and operational management to make servers cooperate (or at least not interfere) with each other.

Vendors and users of NAS devices often assert that the networks connecting their devices to host computers are SANs, and they are right. According to the SNIA, a storage area network is any network that is predominantly used to transfer data between storage devices and computers or other storage devices. NAS devices are certainly storage devices and the mix of traffic on the networks that connect them to their clients is certainly dominated by storage I/O in any reasonably busy information processing operation. We mention this to only nip in the bud any possible confusion around this use of the term SAN. One can only hope that over time the vendor and user communities will evolve to a broader and more descriptive common understanding of the term SAN.

Of course, it's reasonable to ask, What about the other possibilities—block-access devices that use Ethernet to attach to clients and file servers that use Fibre Channel? Both of these are possible and, in fact, development is occurring within the storage industry that may result in both types of devices being brought to market. Today, however, almost all NAS devices are Ethernet-attached file servers and almost all SAN devices are Fibre Channel-attached block-access devices. Table 1.1 summarizes the state of affairs in SAN and NAS devices today.

Summary

Table 1-1. SAN and NAS Storage Device Taxonomy

		INTERCONNECT AND PROTOCOL	
		Ethernet and TCP/IP	Fibre Channel
FILE SYSTEM LOCATION	File System in the Client	"iSCSI" protocol being standardized; products available today	Today's SAN devices
	File System in the Storage Device	Today's NAS devices	Hardware and software components exist, but products not widely available

- A storage area network (SAN) is any network whose primary purpose is to enable storage devices to communicate with computer systems and with each other.
- The key feature that defines a SAN is any-to-any connectivity of computers and storage devices. SANs reduce or eliminate the incidence of unconnected islands of information.
- SANs can reduce the cost of providing information services. In some cases, they can enable new services that it was not previously practical to provide.
- At the current state of this new industry segment, SANs' potential has barely been touched. Ultimately, SANs may lead to heterogeneous distributed systems in which applications running on different operating systems and hardware platforms can meaningfully access the same data.
- Already today, SANs are enabling new ways of providing information services, including an entire new industry segment called storage solution providers, who make quality of service-based access to storage available to their clients as a service.
- Strictly speaking, NAS devices connect to SANs, although the two acronyms are most often used with mutually exclusive meanings. The key unique feature of NAS devices is that they place the file system on the storage device side of the network. Today, the only widespread incidence of heterogeneous data sharing is with NAS devices.