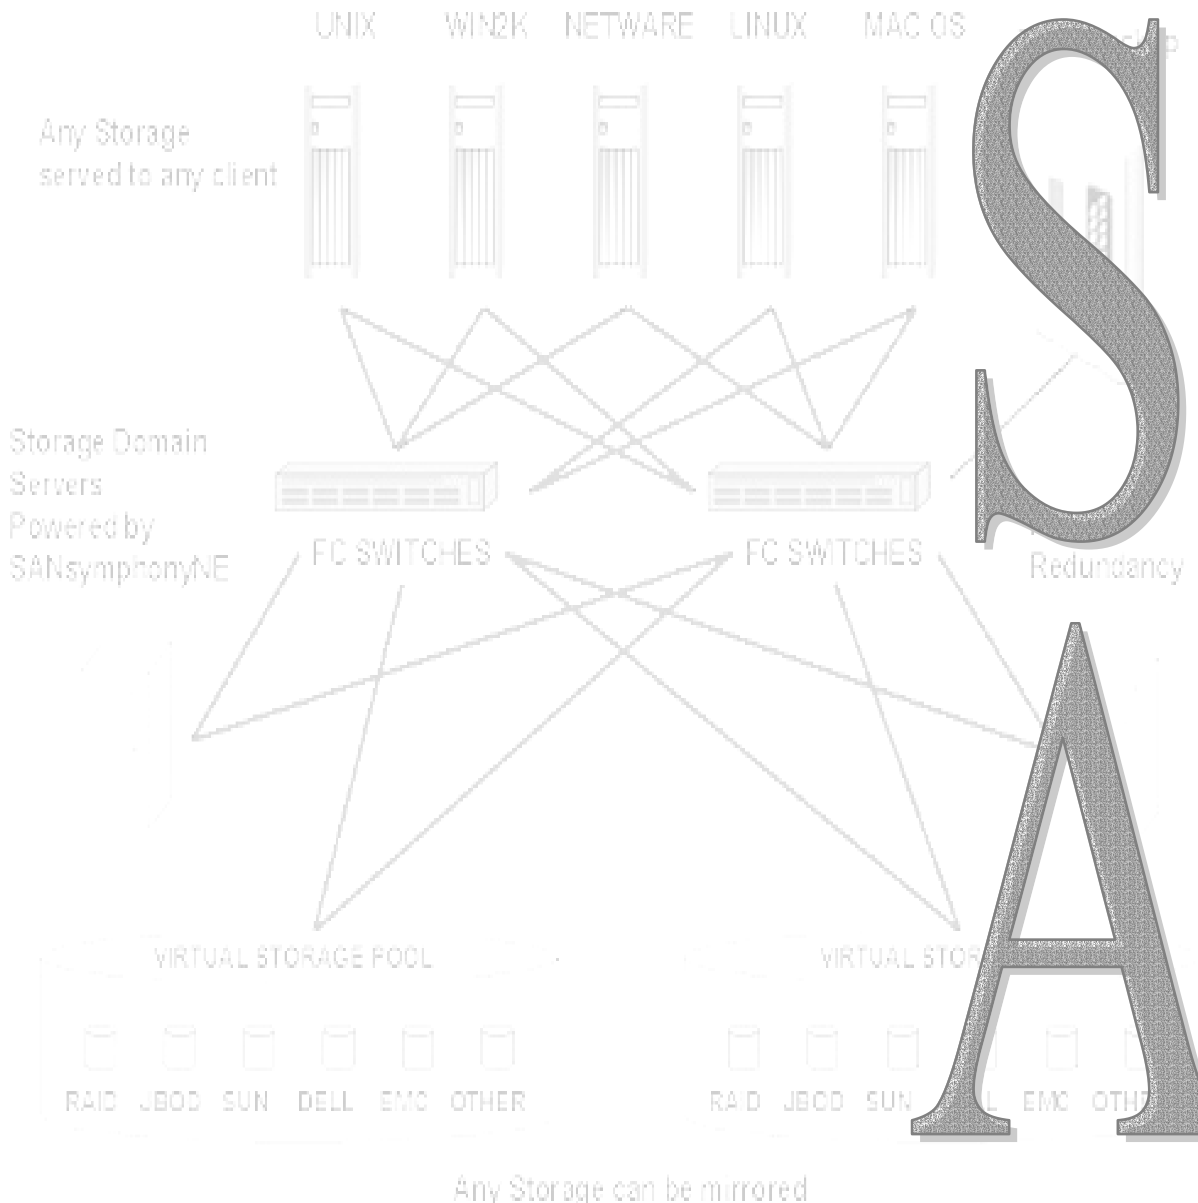

SAN

Storage Area Networks

Niharika Kothari

April 24, 2003

Index

Abstract

This paper was written with the intention of providing information on Storage Area Networks (SAN) and justifying its efficiency. The paper discusses the architecture of SAN and its usage in large enterprises. It also describes in different topologies of SAN. The need for SAN and management of SAN is also discussed. It also sheds some light on why we need SAN and why it is better than most other storage devices and technologies that we use today. The security features and issues of SAN are also discussed. At the end, a few latest developments of SAN and the future of SAN has been discussed. This paper also features the development of Fibre Channel technology for SAN implementation.

Introduction

Technology has come a long way since the man first made the first computer. Ever since, the demand for better and faster has been on the rise. Technology has brought the computer from the 5-storeyed building onto our palms. The 90s saw a huge rise and demand for the personal PCs in the world market. Along with the demand for better machines, the market for storage devices has also been growing. This is mainly because of the growing size of the enterprise and transactions that take place everyday. Storage of information is a very critical issue and has been one since the very beginning. We have come a long way since the first storage devices that hit the market. Today, with growing technology, optical storage devices have hit the markets and are doing very well. All this is very useful for the growing world market. The problem of storage might be solved for small enterprises but not for the big ones that deal with data every second of the day. Here comes the need for storage devices that can store huge amounts of data and also provide disaster recovery features. Huge disk arrays and tape drives can serve the purpose but has a physical limit to how big they can be. Also, they are usually directly connected to the servers, which might not provide disaster recovery. In this case, the back up is usually stored in a different location that is not connected to the server or for that matter even to the network. One could find a whole building filled with back-up files. These operations would also require a lot of man-hours and data might not be available when required, since the back-up files are at a different location. This is where the need for Storage Area Network sprang up. These networks interconnect different storage mediums and offer very high-speed transactions. One of the unique features that it offers is of data recovery. It automatically updates the back-up files every time a new document is updated on the server. Another important feature that it offers is that it is not directly connected to the server and hence offers a very unique security feature. In the pages that follow, a detail study of the SAN architecture and features has been discussed. This technology is the future of storage.

What is a Storage Area Network (SAN)?

In simple words, a SAN is a storage network, which offers large amounts of data storage. These storage devices range from tape drive to disk arrays. It is designed for enterprises that require huge volume of data storage. It is a high-speed network architecture, which is connected to the servers and computers with the help of Fibre Channel hubs, switches and bridges. In this architecture of network storage, the computers and servers are not attached directly to the storage devices.

There are a lot of misconceptions about what a SAN really is. A SAN is not embedded storage. In embedded storage systems, the disk drives are located in the server. These may be high-storage devices, but there is still a physical limit to its storage capacity. Embedded systems are good for small applications. Their biggest disadvantage is in the case of a power or system failure, a lot of data will be lost. SANs on the other hand, are scalable.

A SAN is not directly attached storage either. In this case, all the disk arrays are directly connected to one server. This is also not a SAN. The limitation of this type of storage architecture is the number of bus adapters and addresses available to the server.

A SAN is not Network attached storage (NAS). NAS is very efficient in various business operations. It is very easy to add extra storage to this network of storage devices. It also has a back up tape drive. But it yet does not qualify as a SAN. NAS is directly attached to the Local Area Network (LAN). This makes it less efficient than SAN since attaching it directly to the LAN can degrade its performance. Hence, for larger operations, this form of storage architecture becomes inefficient.

In a NAS environment, both the client and the server request for the records over the LAN. Due to this, at times there is a huge build-up of traffic on the LAN. Also, NAS puts storage where the I/Os impact the clients.

A SAN can be distinguished from other architectures since its storage lives behind the server and not on the LAN. Also, multiple servers can share multiple storage devices in a SAN environment. A typical SAN environment is shown in the figure below.

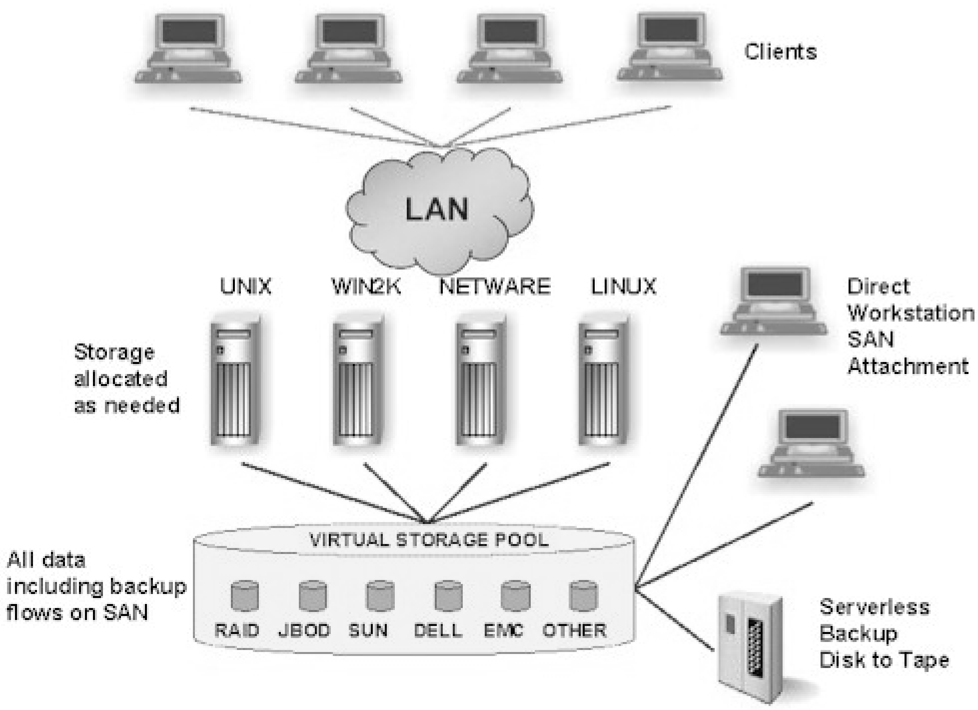

Figure 1: A Typical SAN Environment[1]

Some of the mass storage devices used in a SAN environment are tape libraries, disk arrays, JBODs and optical storage devices. These devices are accessible to all the servers through hubs, switches and bridges.

[1] Picture Taken from http://www.infrastor.com/tech/ SANTechIntro.htm

SAN Storage Devices

One of the main features of SAN is that it can store huge amounts of data with ample opportunity for expansion. The number of devices is limited to the number of interconnecting servers and hubs. The storage devices used are disk arrays, JBODs, tape libraries and other optical storage devices.

Disk Arrays are one of the extensively used storage devices in a SAN. It can store up to terabytes of data, depending on the size of the array. Also, multiple disk arrays can be interconnected to support the needs of an enterprise. Also, it is very fast and highly redundant. It also supports multiple paths to the SAN, which helps in eliminating a single point of failure.

Just a Bunch of Disks or JBODs is a collection of disk drives, which may or may not have high availability characteristics. JBODs can be connected to the SAN with through Fibre Channel bridges.

Tape Libraries are another kind of storage technique that is found in a SAN environment. Tape libraries of any scale can be a part of SAN. Multiple tape libraries are very efficient in storing back-up copies since it is fast and reliable.

In a SAN, it is possible to have direct communication between storage devices without server intervention. Hence, disk-to-disk and disk-to-tape backups are possible. This leads to server-less back-ups.

SAN Interconnection Devices

SAN uses devices such as hubs, bridges and switches to connect with each other. The medium of connection is Fibre Channel.

Hubs are widely used in SAN to increase distance between devices and the number of ports available for connecting devices. One limitation of using a hub is the physical limit of the number of ports available in a hub. Another problem might arise in the performance due to the limited bandwidth. If there are too many devices attached at the same hub, the response time may be slower than usual.

Switches are now gaining prominence and replacing hubs in some situations. Switches have more connection ports than hubs. Though it is very expensive when compared to hubs, they are more efficient since they can carry out more processes at once than a hub.

Bridges are essential to bring Small Computer System Interface (SCSI) on to a SAN. The bridge is necessary to connect SCSI equipment to a SAN. Also, it is necessary to for tape backup, since no tape libraries offered at this time in the market are SAN ready.

SAN Connection Medium

Fibre Channel is the most widely and efficient way of connection for a SAN. This is because Fibre Channel is very fast and can cover long distances. At present, Fibre Channel moves data at gigabit speed. This makes it a very efficient tool to be present in a SAN environment since high-speed data transfer is required. The main purpose of using a Fibre Channel is to connect the storage devices and servers over a SAN.

Fibre Channel can connect devices over long distances. There can be distance of ten kilometers between hubs and the connection is possible through Fibre Channels. Asynchronous Transfer Mode (ATM) can also be seen in a SAN. With the help of this technology, it is possible to move data at will over continents. SCSI connection generally is visible when Fibre Channel is absent. They can be connected through bridges and the bridges are connected to the SAN with Fibre Channel. Different technologies play different roles in connecting a SAN. Each has its own use and works efficiently in its place.

The main drawbacks of networks today are speed, distance and number of devices supported.

Fibre channel storage solutions have destroyed distance, capacity, connectivity and bandwidth limitation in a network environment. It allows high-performance, highly scalable storage solutions for the most challenging situations. Fibre channel has proven to be one of the best interconnecting technologies for network storage solutions.

Fibre Channel uses lasers to achieve high data transmission rates. Turning a laser on and off generates the ones and zeros; this is used to convey information digitally. Receivers detect these bursts of lights and convert them into electrical signals. Fiber has significant bandwidth, transmits data over long distances, and transports data more reliably than copper cable. Also, it does not radiate energy and is immune to electronic noise.

Fibre channel networks provide enterprises with new levels of performance and reliability. Fibre Channel was developed by the industry for IT applications.

Fibre Channel Topologies

The simplest of the Fibre channel topologies is the point-to-point topology. It simply connects two devices with a single link.

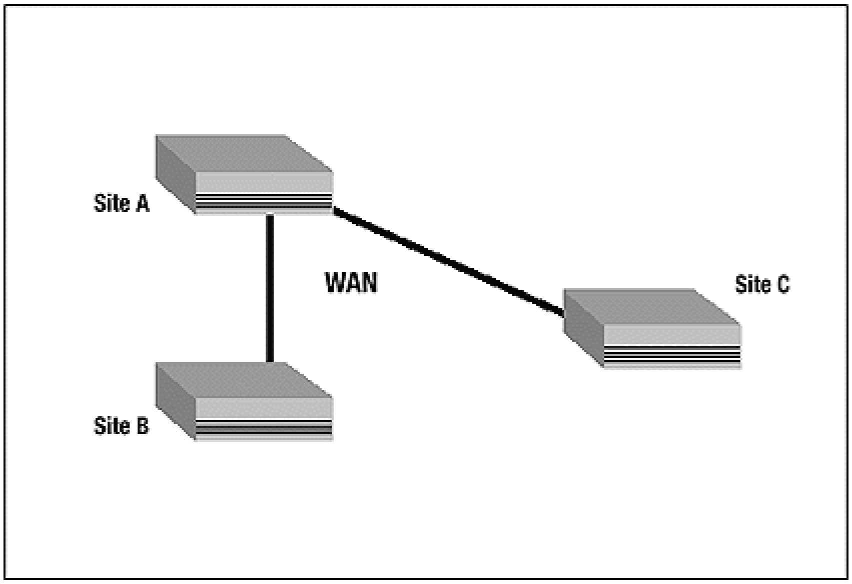

Figure 2: A Simple Point-to-Point Connection[2]

The next one is the arbitrated loop. This topology allows up to 127 ports connected in series, continuing around the back to the origin. An example of this topology would be a network where the first node is connected to the second, the second to the third and so on till the last one is connected to the first, hence completing the loop. But even with this topology, at the time of the connection, it is a point-to-point connection or a two-node loop. In this loop, only the two connected loops can communicate at a given time. Once the transaction between two nodes is completed, a new connection can be established.

[2] Diagram taken from http://www.intel.com/.../white_papers/ pix/widearea_fig4.gif

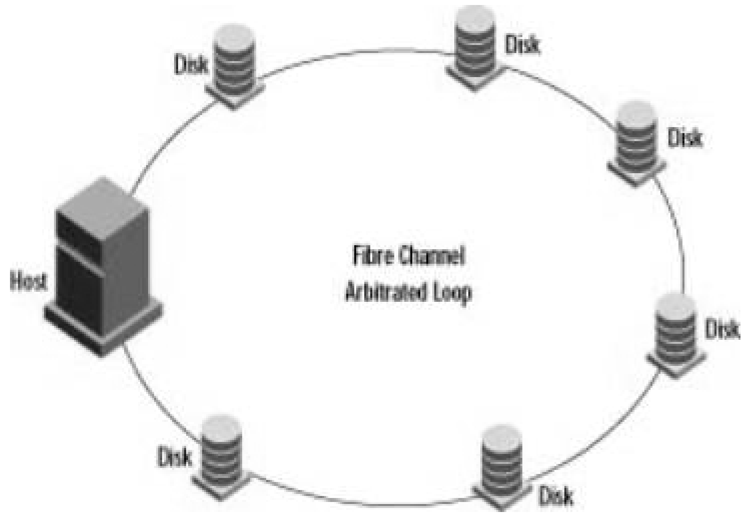

Figure 3: A model of the Arbitrated Loop[3]

The third topology is the switch topology. This topology uses the concept of the fabric, which is a mesh of connections. When attached to a fabric, a single port can access all the different ports and loops on the fabric. Due to this, the number of connections can be increased as compared to the other two topologies.

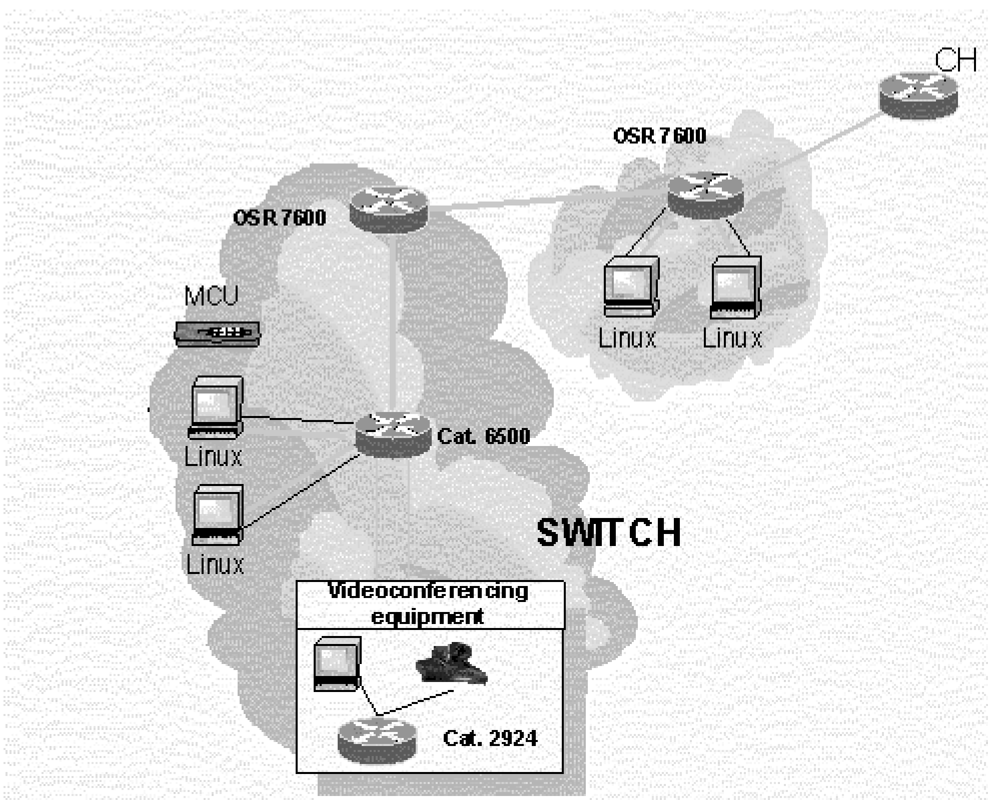

Figure 4: The Switch Topology[4]

[3] Picture taken from http://www.syngress.com/book_catalog/ 140_san_broc/chapter_01.htm
[4] Picture taken from http://www.dante.net/sequin/workshop/ wp5/sld011.htm

Fibre Channel Functional Levels and Protocols

The Fibre Channel functional levels can be divided into two parts – Node level and port level as shown in the figure below.

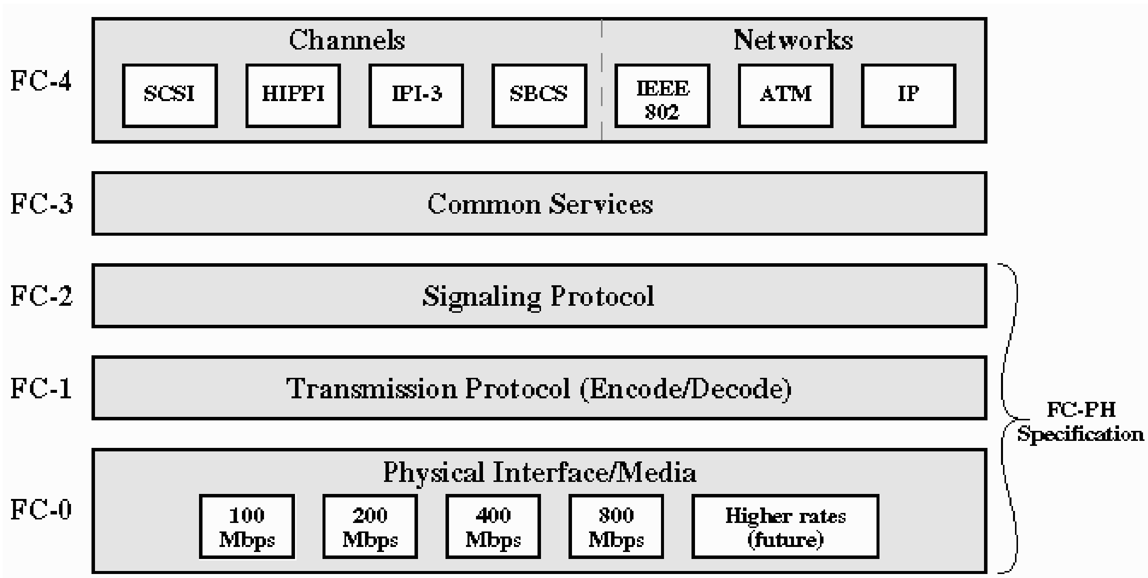

Figure 5: The Functional Levels of Fibre Channel[5]

FC-0, FC-1 and FC-2 are implemented at the port level whereas FC-3 and FC-4 are implemented at the node level. Also, a link that is not visible on the channel is considered system dependent, and the functions are automatically identified. It does not require any allocations or placements. The diagram below shows the framing levels of the Fibre Channel.

[5] Picture taken from http://users.uma.maine.edu/faculty/ rsm/slides/figures.htm

SAN Servers

A SAN implements multiple server environments. A SAN pools the data between multiple servers with relative ease. This lessens the time dependence on any one server and makes the architecture more efficient. Also, it is easy to add more storage without any hassle or loss of data. The servers get instant access to the newly attached storage. Since they are not inter-dependent, servers can be added too. Using this architecture lessens maintenance costs. The servers can be a combination of heterogeneous group. The SAN allows the use of multiple and various server environment.

Brief History of SAN

The SAN architecture is not an innovation that happened overnight. It has taken a lot of time and manpower to develop this technology. There are a lot of reasons for this network storage to emerge and become successful, first of all, the need for large and fast storage devices. This is not a new field, but has been around ever since the first computer was invented. Different techniques have been tried before to provide enough storage for an enterprise, but the problem is it is never enough. Let us look at the different technologies and idea that have led to the innovation of SAN.

Tapes

Tape drives have long been the most widely used form of storage media. There are several reasons for this. It is very easy to use and available readily at cheap rates. The architecture is also not very complicated. The computer memory reads and writes data from and to the tape drive, which is directly connected to the server. But the major problem with this type medium is that its storage capacity is limited. Immediate back up of files is also not possible with this type of media. Also, it requires a lot of manpower and is not the best way to secure data. Tape drives offer sequential access to data that makes it really slow. All these short coming led to the development of disks.

Disks

This form of storage media is also very efficient and one of the mort widely used media even today. It offers direct access to the files that makes it much faster than tapes. But it also has its shortcomings. There is a physical limit for the size of the disks. Also, storage cannot be added to the same disk. A new disk has to be used in each case. Immediate back up is also not available. It is directly connected to the server and the users, which made it less secure. Also, the users could not share data at the same time. These shortcomings of disks and tapes led to the first network access idea, LAN.

Local Area Network (LAN)

This was one of the most important innovations of the twentieth century. This technology solved a lot of problems. Multiple accesses were made possible in a much secure environment. Also, storage problems lessened since LAN connected the different servers to one another. It was much faster than the previously used storage. Also, data was available when required. This architectures led the engineers to think of more sophisticated ways to store data. This led to the invention of networked data storage systems. There was still a problem; the security of data and access was still an issue that had to be solved. The next kind of storage architecture that came up very recently was Network Attached Storage (NAS).

Network Attached Storage (NAS)

This architecture is relatively new and widely used. This architecture attaches the servers, clients and the storage devices all on one network. This is very efficient to store huge amounts of data, but yet has a shortcoming. The storage is in front of the server. This means that the data is available to any ser on the network; this is a security issue again.

Mass Storage

This architecture simply connects the storage devices to one another to provide huge storage space. These disks are independent of each other. This is one of the widely used architectures even today. But this storage medium also has its shortcomings. It is comparatively slow for such a huge environment. The distance between the devices is also limited. It supports a limited number of devices. All these disadvantages led the industry to work on new architectures.

Hence emerged Storage Area Network, which used Fibre Channel to connect with the servers and within each other. This architecture is the very new and will take some time to create an impact in the storage media market.

SAN Principles

Like any other technology, SAN too has its advantages and disadvantages. This architecture of storing facility was initiated by the ever-growing need of the required storage area in large enterprises. Some of the features of SAN include:

- Multiple Storage devices connected to each other
- Storage behind the server
- Multiple server connected to the same storage
- Different types of servers can be connected to the storage area
- Fibre Channel connectivity
- Use of Hubs and Switches
- Multiple path environment

These features might not be available on all the SANs. It depends on what type of architecture the enterprise chooses. The SAN might not have multiple paths to devices on the network. It might not serve the purpose of serving several heterogeneous servers. There are certain risks involved when choosing SAN as the storage architecture.

SAN Building Blocks

The SAN architecture looks similar to the once shown in the figure below.

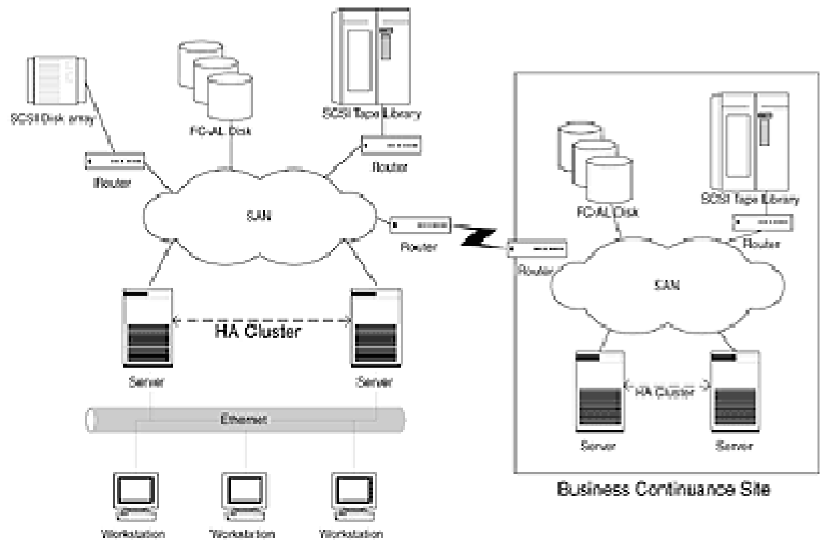

Figure 6: SAN Building Blocks[6]

A whole network of workstations is connected to the corresponding server. The servers can be different from one another. These servers are not attached directly to the storage library, but with the help of hubs and switches. These hubs and switches connect to the storage disks, which again can be multiple in numbers. The medium of connection in this network is Fibre Channel due to its high-speed and bandwidth. Fibre Channel has been a boon in the storage network topology.

[6] Picture taken from http://www.unisysworld.com/monthly/ 2000/07/stornet.shtml

Failure Proofing

RAID Technology

RAID stands for Random Array of Independent Disks. This technology groups the disk drives into a single disk unit and functions as one or more virtual disks. This in turn improves the reliability and performance of the network. RAID uses parity data to allow the disk array to function even after a disk failure. Once the failed disk is replaced, RAID helps to automatically update the lost data from the back-up devices without creating any hindrances for the users. This is one of the features, which are highly recommended in such an environment for reliability purposes. The industry has defined different RAID levels. These are:

- RAID-0
- RAID-1
- RAID-1/0
- RAID-3
- RAID-5

Before discussing the above levels in greater detail, the techniques of disk-striping, parity, mirroring and global hot spares have to be understood.

Disk Striping

This is a technique where data is written to and read from uniformly sized segments across all the disk drives in a group simultaneously and independently. The uniformly sized segments are called block stripes. By allowing multiple sets of read/write heads to work on the same I/O operation at the same time, disk striping can enhance performance. The smaller the stripe element size, the more efficient is the distribution of data written or read across the stripes on the disks in the group. The amount of data written or read from each drive is called the stripe element size.

Mirroring

Mirroring maintains a duplicate image of the data in another disk array. The copy is called a disk mirror. This is useful in case of a disk failure. The data can be accessed from the mirror image and can also be recovered once the failed disk is back into function.

Parity

This is a data protection feature that makes data highly available. Parity helps groups to survive failures without losing data. If one disk fails, the other disks can recreate the data from the remaining data and parity information. If the parity disk fails, it can be reconstructed from the other data disks. Parity can be calculated by using a serial binary exclusive OR of the data segments in the stripe written to the disks. The group can survive a single internal bus failure if each disk in the group is bound on a separate internal bus.

Global Hot Spares

This is a dedicated, online, back-up disk used by a disk array as an automatic replacement disk when a disk in a group fails. This also increases the availability of the disk arrays in the group. When any disk in a group fails, the disk array controller automatically begins rebuilding and recovering the failed disk on the global hot spare. Once the disk has been recovered, it starts copying data from the global hot spare onto the new disk. This is most useful when the highest data availability is a requirement. It reduces the risk of a second disk failure. Also, the need for an array manager can be eliminated since the failure can be automatically detected. Multiple global hot spares can be configured in the same environment where data availability is a crucial issue. Thanks to global hot spares, data remains accessible even if multiple disks fail from the same group.

The above-discussed technologies are implemented in the following RAID levels.

RAID-0 (Non-redundant Individual Access Array)

This level of RAID connects about 3-16 disk units. It uses block-striping technique. It helps spread the data across the different disk units, and allows access to simultaneous Input/Output devices to multiple disks. This level of RAID does not provide ant data redundancy, error recovery, or other high availability features. It is generally not used in large enterprises for high storage networks.

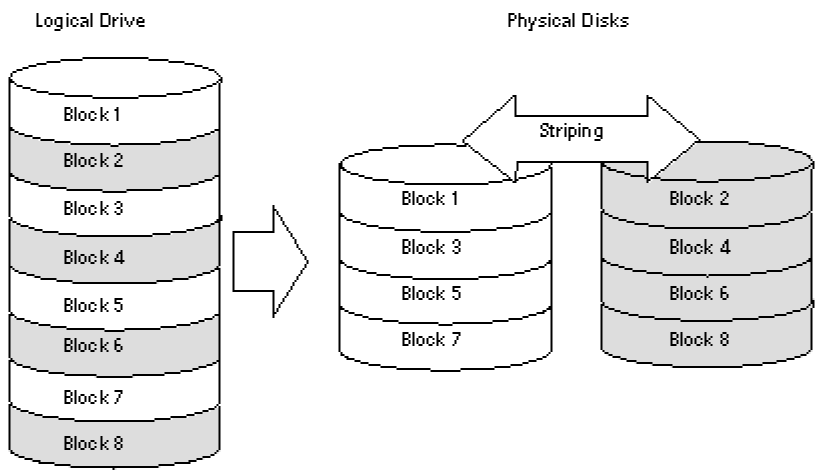

Figure 7: RAID-0[7]

RAID-1 (Mirrored Pair)

This level of RAID contains two disk units bound together as a mirrored pair; one of the disks is the mirror image of the other. The disk hardware automatically writes the data to both the data disk and the mirror disk. Unlike RAID-0, disk striping and parity are not used. In this architecture, if the disk fails, it uses the mirror image to recover the lost data. Just in case both the disks fail, the mirrored pair becomes inaccessible. This is a good solution to save data but requires double the number of disks.

[7] Picture taken from http://docs.sun.com/source/816-7296-11/ ch01_basics.html

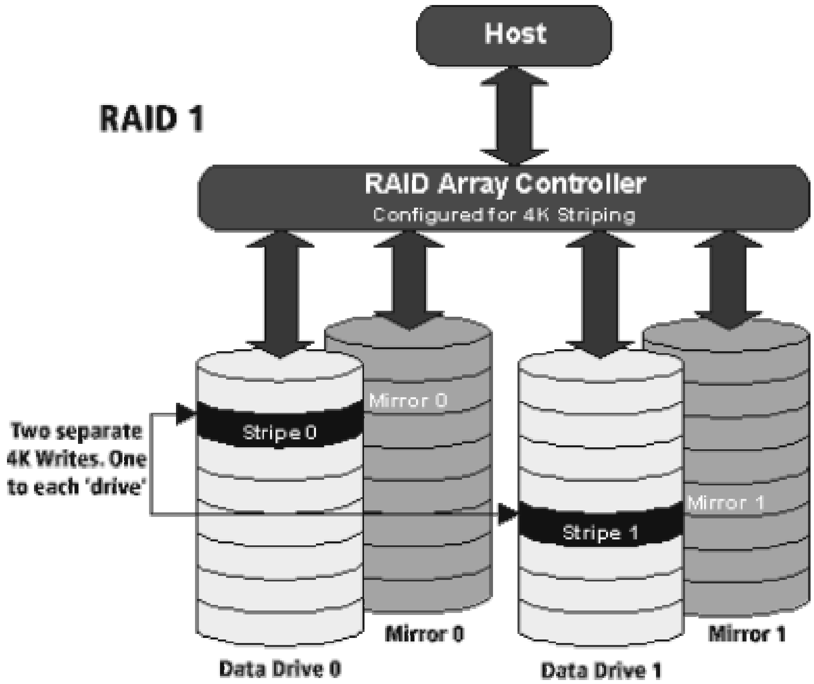

Figure 8: RAID-1[8]

RAID-1/0 (Mirrored RAID-0 Group)

This group is very similar to RAID-0 group. It uses the same topology but implements the mirror-imaging feature of RAID-1. It separates the disks into two sections, one to hold the data and the other to store the back-up data. It still uses block-striping technique for performance and hardware mirroring. The disadvantage of using this technique is that it increases the overhead involved in mirroring as compared to RAID-0 and also the disk space available is reduced. It also requires double the amount of disks, since the data is mirrored. RAID-1/0 is useful in recovering data once a disk fails, since it can recover the lost data from the back-up files. It can also survive multiple failures at once provided at least one of the images is in working condition.

[8] Picture taken from http://www.vogon-data-recovery.com/ disk_recovery-04.htm

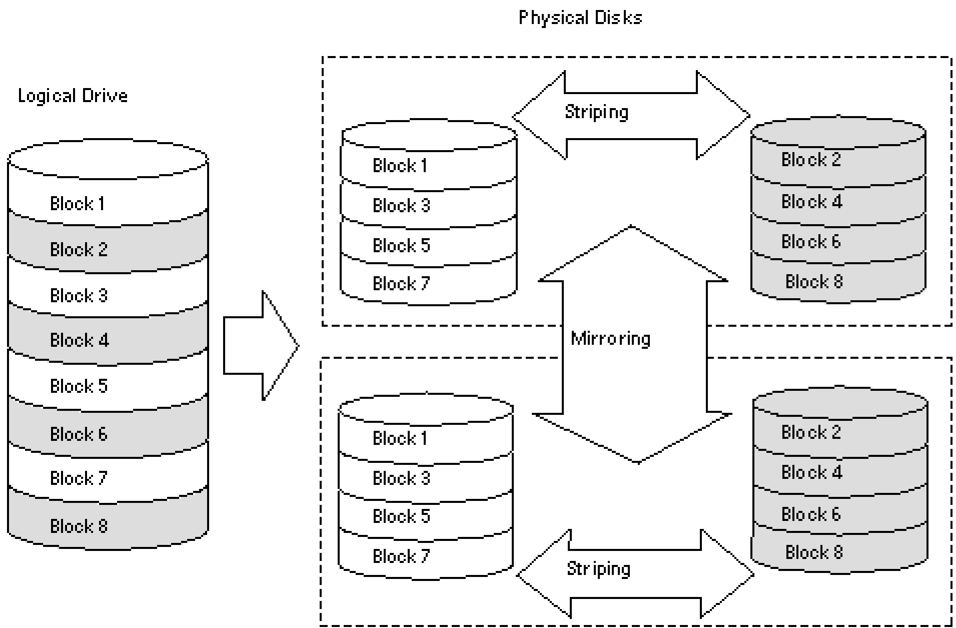

Figure 9: RAID-1/0[9]

RAID-3 (Parallel Access Array)

This group contains 5 disk drives in a disk array. It uses disk striping over 4 disks for performance and the fifth one is a dedicated parity for redundancy. When a disk fails, it can be automatically built again with the available data from the other data. This can slow the performance of the network, but the data is recovered. If 2 of the 5 disks fail, the group becomes inaccessible. This application of RAID is good for large blocks of Input/Output devices.

[9] Picture taken from http://docs.sun.com/source/816-7296-11/ ch01_basics.html

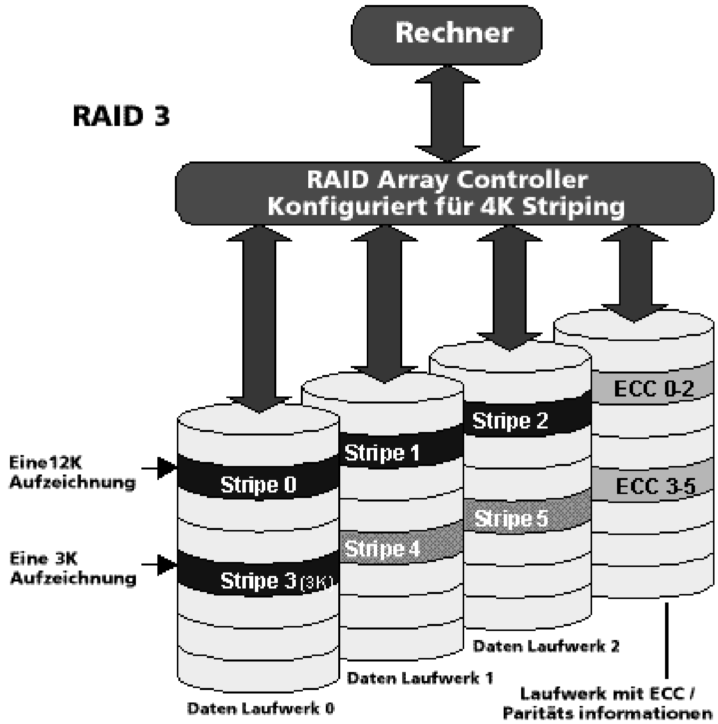

Figure 10: RAID-3[10]

RAID-5 (Individual Access Array)

This level of RAID contains about 3 to 16 drives. Just like RAID-3, this level also uses disk striping and parity; the only difference is that it does not use a separate parity disk. In RAID-5, the hardware reads and writes the parity information on each drive in the RAID group. If the disk array fails, the drive's internal storage processor reconstructs user data and parity information from the information on the remaining drives. This might hinder the performance of the network, but it continues to function with full access to the files. If 2 or more disks fail, the group becomes temporarily inaccessible. This technology is good for multi-tasking environments. It has about the same overheads as RAID-3, but allows faster access since the parity is spread across all the drives in a group. Data transfer can be slower than RAID-3, but can be up helped if the array has caching facilities. It also supports global hot spare technology.

[10] Picture taken from http://www.vogon-data-recovery.com/ disk_recovery-04.htm

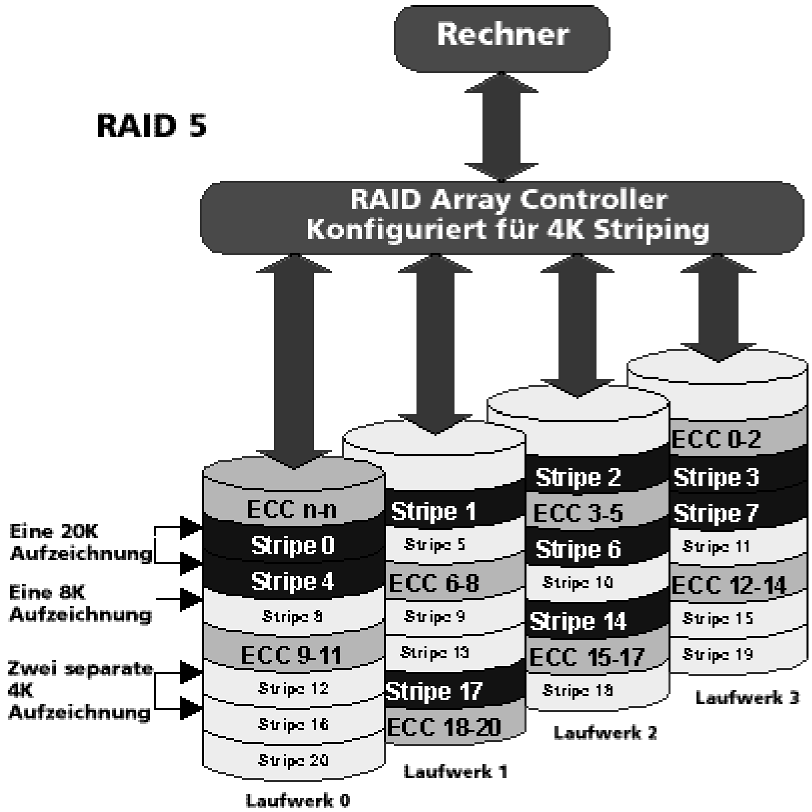

Figure 11: RAID-5[11]

Since RAID takes care of the data redundancy problems, the other major problems that a SAN is susceptible to are controller, fan and power failures. The data can be restored with the help of RAID technologies, but parts cannot be replaced automatically. These components need monitoring on an everyday basis. Hubs shut down failed ports, and thanks to the hub management software this problem can be fixed easily. Defective Fiber optic cable can also be replaced without much hindrance. If the fiber optic cable is bent or defective, it can cause problems in data transfer. But if the cabling is well handled, these problems can be avoided. One of the major advantages of SAN is that in case of a part failure, some other component takes over till the defective part is replaced. One of the major problems that enterprises face is power failure. The best way to guard against this problem is to have back-up power supply available at all times. In times of natural disasters, data can be lost. Hence, it is advisable to have a mirror of the files at a different location. This helps is recovering data in times of disaster.

[11] Picture taken from http://www.vogon-data-recovery.com/ disk_recovery-04.htm

SAN Topologies

As we know, SAN is a network of storage devices that can connect with several servers and workstations through the channel of hubs and switches, which are connected with Fibre Channel. The following topics give a detail study of the topologies of SAN and why it is recommended for use.

Scalability

SAN can be implemented in a large network of hundreds of workstations and several servers. The main advantage of this feature is that servers and storage devices can be added when required. This is very important in a large enterprise where data is stored and used every second of the day. This feature helps in updating the network as the demand arises.

Distance

Fibre Channel technology has proved as one of the greatest boon to the storage network architecture. SANs can now stretch up to several kilometers and still prove to be fast. This technology has helped to store back-up data at different locations and still be attached to the network. The mirroring can be done simultaneously without the need of administrators. This helps to recover data in times of disaster. It also helps in saving expensive realty in huge cities by substituting it with cheaper countryside areas for back-up purposes.

Device Mix

Another important feature that SAN has to offer is that it can combine different servers and storage devices into a single network. Different servers such as Unix, Solaris and others can be attached to the same storage network. Also, different storage devices, such as Disk Arrays, Disk drives, JBODs and other optical storage devices can be implemented under a single network.

Capacity

This is one of the most important feature and one of the main reasons why SAN was implemented – capacity. This is one of the main concerns in the storage world today – how to increase the capacity without disturbing the network. SAN is the answer. Multiple storage devices can be attached to each other to offer large amounts of storage. Also, more storage can be attached later as and when the need arises. This is one of the main reasons why SAN is rising in popularity charts compared to other techniques.

Back-up Facilities

Another reason why SAN is gaining importance is because of its back-up facilities. Data can be mirrored automatically and simultaneously without any human intervention. This reduces the overheads associated with data back up. Also, data can be recovered in case of a disk failure without turning off the network and creating obstacles. Fibre channel makes this possible due to its high transfer rates and distance features.

Implementing a SAN

There are certain steps that need to be taken in order to implement the SAN successfully. First of all, planning is very important. If planned well in advance, one can implement the SAN with fewer problems and obtain better results. Some of the questions one has to answer are:

- What kind of network is planned?
- How large should the storage be?
- How many workstations and servers are going to be implemented on the network?
- How much should the cost be?
- How fast will the need for additional storage grow?
- What will be the future needs?
- What should the SAN look like?
- Where will the equipment be stored?
- Where will the back-up equipment be stored?
- How will the SAN be monitored and managed?

If the blueprint of the SAN network is good, one should have no problems implementing a SAN environment. Implementing a SAN is not very difficult. With the latest technology and available equipment, SAN networks are very easy and efficient to install and manage.

SAN Management

SAN is easy to maintain as compared to the earlier network storage technologies. Depending on the type and size of SAN, the need for management can be eliminated. Some of the tasks that a SAN manager should perform are:

- Monitoring the devices on the network
- Monitoring the functionality of the devices
- Determining if any device needs changing or upgrading
- Enable or disable hubs and ports
- Check the error status and replace the corrupted devices
- Look into the performance of the network
- Controlling the configuration of Disk Arrays
- Setting levels of security for each user
- Allocation of assets within the network
- Performing preventive maintenance

These tasks are not very tedious. The SAN does most of the error reporting automatically. Due to the efficiency of SAN, overheads costs have been reduced. Thus, managing and maintaining a SAN is an easy job.

Back-up

Duplicating data can be a tedious and often thankless job. But even then data backup is required almost everywhere. There are firms that indulge in data backup on an everyday basis. This is important for many reasons. First of all, the data can be used over the years by hundreds of employees. All the saved work cannot be risked to loss. Here comes the need for efficient backup architecture. SAN has helped greatly in providing this feature.

Backup is an operation where the data is copied and stored on a different storage device, sometimes at a different location. It also describes the collection of copied data. The data is copied from primary storage to secondary storage. The size of storage devices and backup storage devices have decreased over the years, but the principle has not changed much, the data is still duplicated and stored in another location.

The main concern of data backup is the immediate duplication of files. Also, it should be compatible with the different servers and should be able to upgrade when new devices are attached. Traditionally, during backup, the system is unavailable to the users. Hence, most of the backing up is done during the wee hours. But today, new data is added every second of the day. Hence, this system might not work. Also, there are a lot of overheads attached with the traditional way of data backup. There have to be a system administrator and other staff for updating and maintaining the backup system. Hence, the cost of backup increases significantly. Also, restoring lost data can be a tedious job since the devices are not attached to each other. Restoring means coping saved data from secondary to primary storage on a network. Loss of data can also occur if copying is not done on a quick basis. For example, a program saved on Monday can be lost by Wednesday, since it was not updated to the backup files immediately. Implementing SAN can reduce these problems to a great extent.

SAN can solve the problems of backup. SAN offers high-speed, immediate backup solutions to large enterprises. It has proved its efficiency by decreasing the overheads. Also, SAN backup storage devices can be several miles away. It is highly scalable and offers high-speed data transfer. Creating backup facilities in a large enterprise is a big challenge. The backup has to be upgraded and new backup devices have to be added from time to time as the need of the enterprise increases. SAN does not interrupt any processes while duplicating data, nor does it slow down any processes. It is reliable and the customer can be satisfied that the data is getting saved. San also promises backup from several heterogeneous servers. All the different servers can access the data from the same storage area network. SAN implements different storage devices in the same network. Disk Arrays, Tape drives, JBODs and other optical storage devices can be installed on the same network. Another feature that SAN offer is that of Disaster recovery. Data can be restored from the secondary storage devices in case of a system failure in the primary network. This is done on a very high-speed basis and hence does not limit network use for any of the users.

There are a few limitations even today for using SAN backup services. The connectivity is limited. Also, the bandwidth is a concern. There might be degradation in the speed for some processes when data is being copied. Error handling is another concern since it needs to be closely monitored. These limitations can be overcome by using future technologies.

SAN Security

Security has been an issue in large enterprises time and again. This has led to the development of more sophisticated ways of securing data from hackers and eavesdroppers. Security of safe backup is also an issue. The exchange of data in an enterprise network has to be handled efficiently in order to safeguard the data. An enterprise requires an integrated solution that looks into the security threats, enabling a robust SAN environment. Since the data is transferred over several miles and sometimes even over several continents, the security features have to be mission-critical. There are a few domains that can be integrated in the system for better security. These are

- Administrator-to-Security Management Domain: an administrator works with the security management functions.
- Host-to-Switch Domain: Individual ports are grouped and linked to one port.
- Security Management-to-Fabric Domain: This encrypts the appropriate data elements with the switch's public key.
- Switch-to-Switch Domain: Switches are initialized and they exchange credentials during mutual authentication, before establishing any communication with the user.

Implementing SAN makes data highly available to all the users. Another problem arises here; all the data entered cannot be made accessible to all the users. Every user needs to have a separate level of security and data access according to requirements. One way of doing this is by Zoning. Zoning creates barriers on the fabric to prevent groups from interacting with each other. Zoning sets up certain levels of accessibility to the network and hence data is not available to all the users.

Future of SAN

SAN has created a stir in the storage industry. It has proved to be highly efficient and is able to store huge amounts of data. One of the drawbacks of SAN is that the storage takes too much of space. For example, the disks and tape drives are huge in size and cover a lot of real estate. The size can be reduced with the help of Nanotechnology. Decreasing the size has proved to be helpful in the past. Today, one can own better and more efficient computers than ENIAC on the palm of their hands. Also, more storage is available. Technology has moved from tape disks to floppy disks to zip disks – decreasing the size and increasing the storage capacity on the way. In this way, storage capacities might increase without the increase in size of the equipment. The interconnection devices, such as hubs and switches need to allow more ports for connection. The backup facilities have to be more efficient and faster and reliable. The ones present today are good for the enterprises of today. But as the enterprise grows, more and more data will be stored. The number of users will also increase. The use of SAN will increase as it becomes more and more popular. Today, SAN is an expensive solution to use as a storage medium. This is mainly due to the high initial costs of installing and buying the equipment. These costs will go down in the future when this technology becomes more popular and widely used. Management of SAN, though is easy, will be easier in the future. RAID might help to overcome the problem of unavailable data caused by failed disks. Switches might have redundant components. The number of equipment required to build the SAN might reduce with the help of new technologies.

Changes might take place in areas such as the Internet and wireless computing. The major problem for the data available today on the Internet is that it is limited to the server's capacity. This drawback might be eliminated with the help of SAN since more and more storage can be made available and it is scalable. Wireless computing might become easier. Today, the computers are limited by the capacity of the hard disks. This problem might be eliminated when SAN becomes more popular. With the help of SAN and wireless technologies, storage can be made in an external drive, which can be accessed from anywhere. SAN can act as a wireless storage solution.

Conclusion

In the paper, SAN has been discussed in its fullest. SAN has stirred a revolution in the storage networks. It has been proved to be useful in many applications. The features offered by SAN makes it very reliable to implement in large enterprises where data integrity is top priority. Fibre Channel has helped achieve the goals that were not possible without otherwise. SAN is not very extensively used today due to the expense related with it. But the future holds a lot for SAN and the market will improve as the cost goes down. SAN holds a very bright future in the storage industry. SAN can be used in almost every industry one can think of; medicine, music, education, administration, exploration, industry and other areas.

Bibliography

General Reference

Ralph H. Thornburgh, Barry J. Schoenborn. *Storage Area Networks.* Prentice Hall, 2001.

William Stallings. *Data and Computer Communications.* Prentice Hall, 2000

Gibson, Nagle, Courtright, Lanza, Mazaitis, Unangst, Zelenka. *NASD Scalable Storage Systems.* Department of Electrical and Computer Engineering, Carnegia Mellon University.

Gibson, Nagle, Amiri, Butler, Chang, Gobioff, Hardin, Riedel, Rochberg, Zelenka. *A Cost-effective, High-Bandwidth Storage Architecture.* Department of Electrical and Computer Engineering, Carnegia Mellon University, 1998.

Gibson, Nagle, Amiri, Chang, Gobioff, Riedel, Rochberg, Zelenka. *Filesyatems for Network-Attached Secure Disks.* School of Computer Science, Carnegia Mellon University, July 1997.

Lim, Kapoor, Wighe, Du. *Active Disk File System: A Distributed Scalable File System.* Department of Computer Science and Engineering, University of Minnesota, Minneapolis. Valicent Inc. Trust Services Engineering. Server Products Group, Wind Rivers Systems.

White Paper. *Fibre Channel Technology: Understanding Fibre Cabling and Connecting to Building Infrastructure.* Enterprise Computing, Compaq Computer Corporation, March 1998.

Web References

http://www.infostor.com

http://www.i-tech/Storage%20Info%20Center/Fibre%20Channel%20Technology.php

http://www.seagate.com/support/kb/disc/fcp.html

http://www.vixel.com/whyfibre/faq.html

http://www.hermesgroup.com/whitepaper/Fibre%20Channel/Fibre%20Channel.htm