IBM

Storage Area Networks: Tape Future in Fabrics

Matthias Werner, Jonathan Jeans, Barbara Manger, Charles Silvan

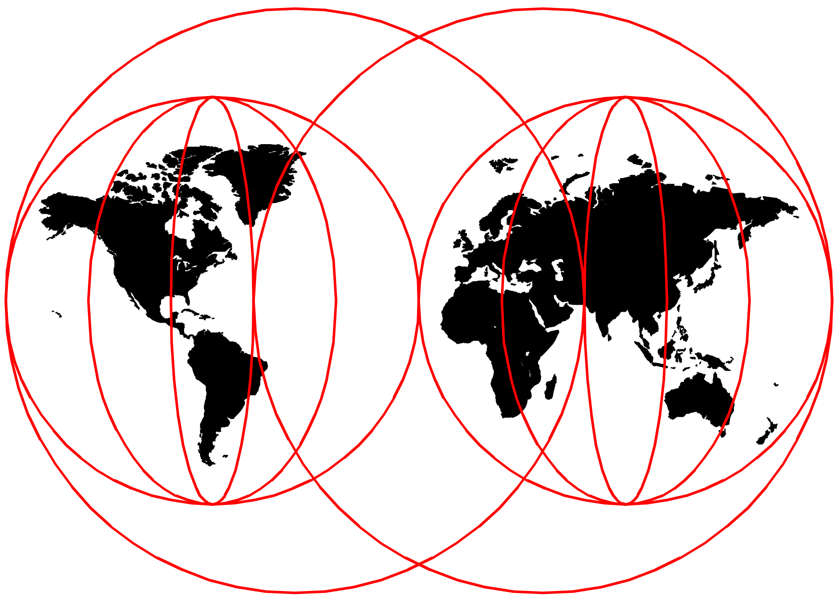

International Technical Support Organization

www.redbooks.ibm.com

SG24-5474-00

SG24-5474-00

International Technical Support Organization

Storage Area Networks: Tape Future in Fabrics

January 2000

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, “Special notices” on page 135.

First Edition (January 2000)

This edition applies to a range of IBM products related to Storage Area Networks, including but not limited to the IBM SAN Data Gateway, the IBM Fibre Channel Switch, the Magstar 3575 Tape Library, and the Tivoli Storage Manager (formerly ADSM).

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. QXXE Building 80-E2
650 Harry Road
San Jose, California 95120-6099

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Figures

Tables

Preface

This redbook explains Storage Area Networks (SAN) and their possible and practical implications and implementations on IBM storage products in general, and tapes and tape libraries specifically. It applies to a range of IBM products related to Storage Area Networks, including but not limited to the IBM SAN Data Gateway, the IBM Fibre Channel Switch, the Magstar 3575 Tape Library, and the Tivoli Storage Manager (formerly ADSM).

This redbook should be read together with a companion redbook, *Introduction to Storage Area Networks, SAN*, SG24-5470.

Written as a follow-on to the *Introduction* redbook, it moves on to the use of SAN technology in tape environments. Tape drives and libraries tend to be expensive devices. The ability to share them across many hosts creates a huge financial advantage that can be one of the immediate benefits of implementing SAN in your enterprise! It includes a practical description of the products and components made available with the first IBM SAN product rollout, and contains IBM redbook-level details on how to make it work.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization San Jose Center:

Figure 1. Almaden Research Center, San Jose

This project was designed and conducted by **Matthias Werner.** He is responsible for tape and tape library projects at the International Technical Support Organization, San Jose Center. His focus is to build and maintain high quality redbooks for use with IBM's leading tape library products in the field. This is a pivotal function between the development labs and customers, Business Partners, and IBM professionals. Matthias also regularly teaches technical workshops and sessions at conferences and seminars. Before joining the ITSO in July 1997, he was involved with most of the IBM tape library projects in Switzerland, gaining broad experience in the tape-centric-computing environment since 1992. He can be reached at wernerm@us.ibm.com.

Jonathan Jeans is a Senior Technical Consultant working for Berkshire Computer Products, an IBM Business Partner. Berkshire specializes in storage management, mostly for Open Systems in the Northeast area of the United States. He has over 30 years of experience in the storage field. He holds a Masters degree in Applied Mathematics from Brown University. His areas of expertise include all areas of storage; tape, disk and storage management software. He actively consults with clients in the Northeast, helping Berkshire’s clients select the proper storage solutions. Jonathan can be reached at *jonathan.jeans@berkshirecp.com*.

Barbara Manger-Strassguetl works for the Advanced Technical Sales Support team for SSD Sales Support EMEA. She has 6 years of experience in the tape arena. Her specialty is Open Systems attached tapes and libraries. For her, SAN breathes life into storage products, thus opening up new and interesting challenges.

Charles Silvan is a Storage Specialist in IBM Global Services in Geneva (Switzerland). Charles holds a degree in physics from the ETHZ and a Masters Degree from Union College, Schenectady, NY. Charles has been with IBM for 29 years, including 3 years at the World Trade Systems Center in Poughkeepsie. He has worked in a variety of areas, including MVS, TPNS, SNA, VM and Document Processing. He is presently working in Services, where he is responsible for storage and ADSM projects in the French-speaking part of Switzerland.

Thanks to the following people for their invaluable contributions to this project:

Kjell Nystrom
International Technical Support Organization, San Jose Center

Lee Jesionowski
Mark Doumas
Donald Denning
Reed Hancock
Karen Orlando
Greg Tevis
Erik Bengds
Michael Kaczmarski
John Kulakowski

IBM Tucson

Scott Drummond
Lisa Haut-Mikkelsen

IBM San Jose

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in "IBM Redbooks evaluation" on page 147 to the fax number shown on the form.
- Use the online evaluation form found at `http://www.redbooks.ibm.com/`
- Send your comments in an internet note to `redbook@us.ibm.com`

Chapter 1. SAN

During the past ten years, a major shift has occurred away from the legacy mainframe host-centric storage model, to the distributed storage that is currently in vogue in the client/server model. Storage, for the most part, remains connected to the processor it services via a dedicated channel, such as a SCSI bus. However, the large number of servers creates a complex environment and is difficult to manage. The high dedicated SCSI channels tend to be inflexible, especially in terms of distance between the storage and the processor. At the same time, in the client/server model, the servers are interconnected via network protocols, such as TCP/IP, so that these servers can communicate with each other. These networks, LANs and WANS, tend to be slow but very flexible. Interconnectivity is easy, but large blocks of data traverse the network rather slowly and network disturbances are frequent.

The industry realizes that a new type of data network is needed to meet the needs for the new distributed environments. The outcome of this process is a set of standards defining Fibre Channel. Fibre Channel brings the best of both worlds together — high speed and robustness of the storage channel with the flexibility and few distance limitations of a standard network. Fibre Channel is based on a network definition of five layers, where the lowest three physical level layers define the pathway and the upper two layers enable the software protocol to move the actual data. See Figure 5 on page 7 for a picture of these layers. (See Appendix A for a complete explanation.)

The uniqueness of Fibre Channel is that these two upper layers can be used to carry older protocols. Protocols such as SCSI, ESCON, and IP can be mapped to the Fibre Channel transport service level. Therefore, reinvention of TCP/IP, SCSI, ESCON, and other protocols is not necessary. The SCSI protocol, and thus SCSI commands, can be transported across the Fibre Channel infrastructure. A newer SCSI-3 protocol has been developed to allow for some of the special functions available with the Fibre Channel standard. The Fibre Channel transport can use either copper or fiber optic cables, can move 100 MB/sec across 10 km, and allows for an almost unlimited number of devices to be interconnected.

Today many vendors are rushing to the market with new Fibre Channel products and putting them together into networks they call Storage Area Networks (SANs). The promise to resolve many of the industry's major issues (speed, connectivity, availability, reliability, and manageability) is creating quite an energetic market place. Using a SAN can potentially offer the following benefits:

- Improvements in application availability: Storage that is independent of application and accessible through alternate data paths.
- Higher application performance: Delegate storage processing to storage subsystem off-loading server.
- Centralized and consolidated storage: Simpler management, scalability, flexibility, availability.
- Data transfer and vaulting to remote sites: Enables remote copy of data for disaster protection.
- Simplified centralized management: Single image of storage media simplifies management.

IBM has launched a SAN Initiative, thereby leaping onto the band wagon. Today several products are available from IBM to begin the process of creating and exploiting this new technology. However, like all other vendors in this area, IBM has found that the new technology is not so easy to implement and get to work seamlessly. No single vendor today markets all the pieces to put together a complete SAN, and it is unlikely that one ever will. Additionally, no single standard exists at this time, but rather there are a set of loose standards that are being tightened as time passes. Thus, today, the major issue facing SAN technology is the interoperability between the pieces of the SAN.

SANs create new methods of attaching storage to processors. (See Figure 2.) These new methods promise great improvements in both availability and performance. Currently they are most commonly used to connect shared storage arrays to multiple processors, and are used by clustered servers for failover. They can create the opportunity for various servers to share tape drives and tape libraries. A SAN can be used to bypass traditional network bottlenecks. It supports direct, high speed data transfers between servers and storage devices in the following three ways:

- Server to storage: This is the traditional model of interaction with storage devices. The advantage is that the same storage device may be accessed serially or concurrently by multiple servers.
- Server to server: A SAN may be used for high-speed, high-volume communications between servers.
- Storage to storage: An example is a disk device backing up its data to a tape without processor intervention. It could also be remote device mirroring across the SAN.

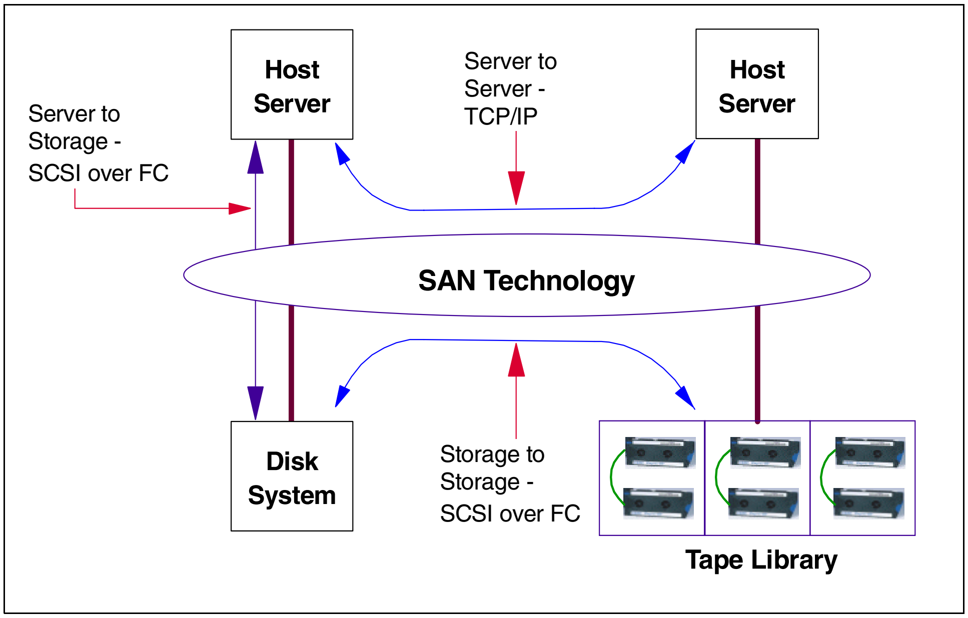

Figure 2. High speed data transfers

We expect that vendors will announce many solutions over the next few years. Many solutions will be effective, and many will lead us down the wrong path. What we would hope is that this book can help you determine your future vision, so you can decide which products you should pursue. IBM has developed many solutions to date, and will continue to deliver solutions at a fast pace over the next few months and years. The goal is to deliver products that solve specific problems, that are economically justified, and that work together with other solutions developed. IBM intends to lead the SAN practice through service expertise which both:

- Allows us to implement specific IBM solutions at a customer site
- Helps customers put together heterogeneous solutions from various vendors and ensures that they work together.

We hope that through this book, the reader will be able to recognize the value of the SAN technology as it relates to tape, and will be able to understand several products available from IBM, to allow exploitation of the SAN technology in the tape environment.

1.1 Definitions and terminology

This section will offer a general discussion of Fibre Channel (FC) architecture and SAN topologies, and should provide an adequate introduction to the terms that we use throughout this redbook. However, to get a better understanding of FC and SAN technology, the reader may also refer to the redbook: *Introduction to Storage Area Network, SAN*, SG24-5470.

1.1.1 What is a SAN?

There are many definitions for a SAN. Most are developed by the various vendors to help market their own products. In this redbook, we will use the following definition, which is quite broad in scope, but requires that some Fibre Channel protocols are installed. Fibre Channel unlocks many limitations placed on the storage environment by the older channel architectures.

A SAN needs to have the following four elements:

1. Some elements of the network, such as Fibre Channel cables, a switch, a hub, or a gateway.
2. At least one host node. (Note that the node resides in the Host Bus Adapter card and not in the host itself.)
3. At least one storage node. (Note that the node resides in the Device Adapter card and not in the storage subsystem itself.)
4. Some management software to manage such things as the devices, allocation, use, and configuration. The main purpose of the software is to enable the storage and host elements connected to the SAN to operate with data integrity.

And here's the IBM definition of a SAN:

A Storage Area Network (SAN) is a dedicated, centrally managed, secure information infrastructure, which enables any-to-any interconnection of servers and storage systems.

Figure 3 shows the main components of a SAN. Thereafter, throughout this document, we will only show a cloud with the words to depict a SAN. The reader should keep in mind that this cloud is a simplification, and that most SANs will have many elements such as switches, hubs, and gateways. However, it is simpler for us to ignore them in the figures.

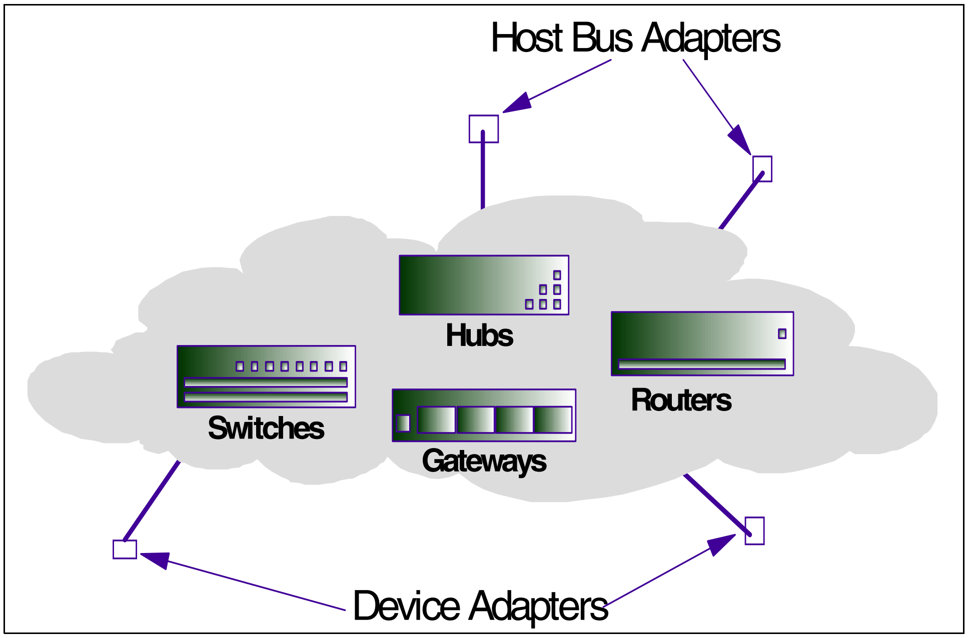

Figure 3. SAN components

These elements, when put together, create a Fibre Channel topology which we will call the SAN. The SAN is nearly limitless in architecture, but most current implementations are very limited in terms of distance and number of connections. The maximum switch size today is 32 ports, while the largest hubs have only 16 ports. (These elements, hubs and switches, and others are defined in Appendix A.) Software management tools are also limited, and tend to be able to manage only a small subset of the SAN. Although these SANs seem immature, they can solve some real-life problems. However, the major reason most of the SANs are installed today is to obtain knowledge and expertise in the technology so that, as it evolves, the enterprise can adopt the new technology quickly and painlessly.

Figure 4 depicts a SAN with several processor types and several storage types. For the remainder of this book we will depict the SAN as a cloud, but the reader should assume that each cloud can be made up of switches, hubs, gateways and bridges, as depicted in Figure 3. Notice that in this depiction the processors are still able to communicate across the LAN. With this level of connectivity comes a burden — and that is, how to protect the integrity of the device and data across the various servers, since each server can actually talk with each device. This is the reason software management tools must be in place to exploit a SAN.

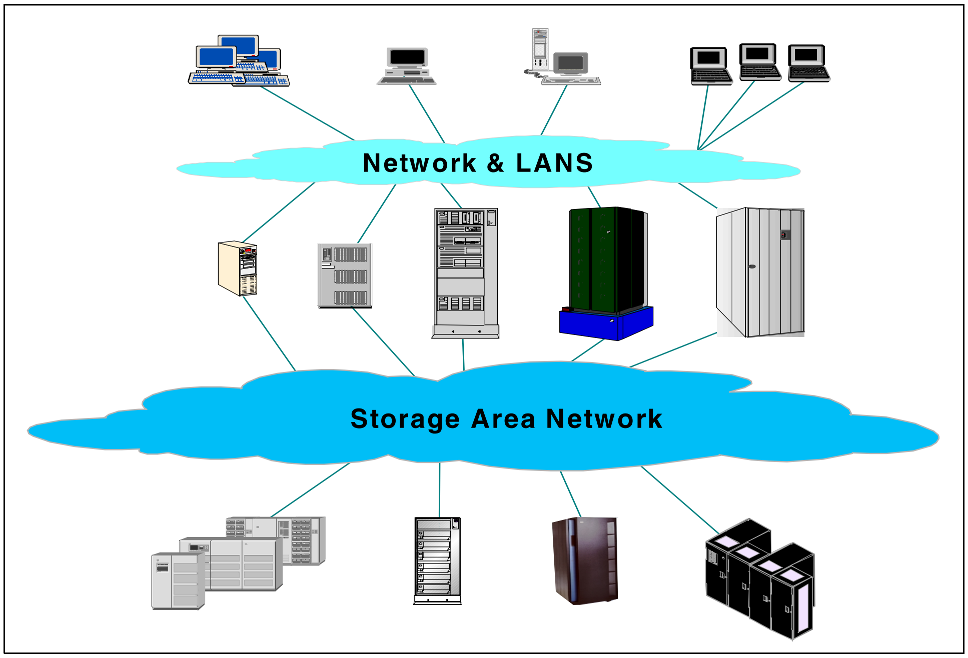

Figure 4. What is a SAN?

1.1.2 Fibre Channel

Today the industry considers Fibre Channel as the underlining architecture of the SAN. Fibre Channel is a technology standard that allows data to be transferred from one network node to another at very high speed. (See Appendix A for more details on the Fibre Channel architecture.) Current implementations transfer data at 100 MB/sec, while 400 MB/sec has already been tested. This standard is backed by a consortium of leading vendors and has been accredited by the American National Standards Institute (ANSI).

Note the spelling of the word *Fibre* in Fibre Channel. This is because the interconnections between nodes are not necessarily based on fiber optics but can be based also on copper cables. In the following sections we will introduce some basic Fibre Channel concepts, starting from the physical layer and going as far as to define the services offered.

Fibre Channel is structured with independent layers, as are other networking protocols. There are five layers, where 0 is the lowest layer. The physical layers are 0 to 2. These layers carry the physical attributes of the network and transport the data created by the higher level protocols, such as SCSI, TCP/IP and ESCON.

As you can see in Figure 5, the two top layers, which are the session and transport layers, can be used by these protocols to move data segments. These segments are then rolled into a packet, which in turn gets rolled into a frame. The source end creates the frame, while the target (destination) unravels the frame back to a segment. SCSI protocol has been extended to SCSI-3 to allow the protocol to travel across the Fibre Channel hardware layers.

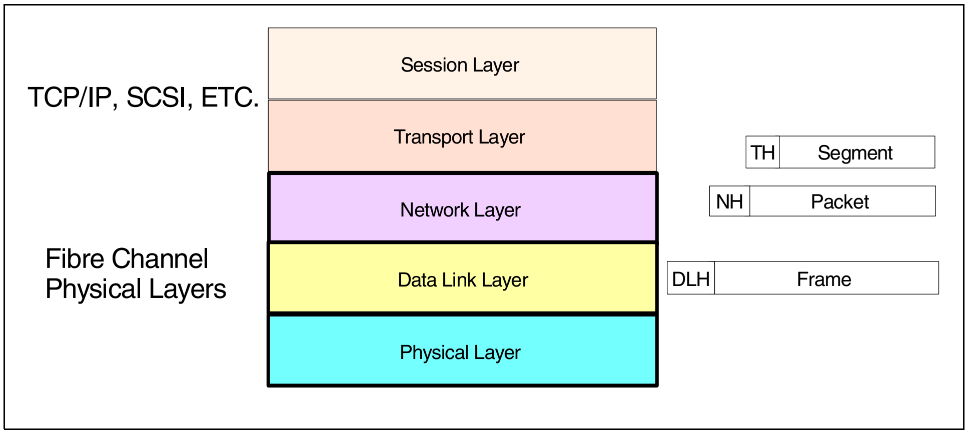

Figure 5. Fibre Channel protocol layers

Fibre Channel provides the physical service to move data fast and with reliability between nodes. The two upper layers enhance the functionality of Fibre Channel and provide common implementations for interoperability. One defined function is multicast to deliver one transmission to multiple destinations.

An analogy with trains and train tracks can be drawn. For example, in the heyday of steam trains, single tracks were set down and steam engines were built. The passenger cars carried passengers. We can think of the passengers as the data to be transported and the track and train as the physical medium to move these passengers (data). Then, we can see that if we upgrade the tracks to welded tracks with high speed jet-like engines to pull the shiny new cars, we have upgraded the service, but the passengers are still people. The storage industry has upgraded the tracks and trains (Fibre Channel) but kept the passengers (SCSI, ESCON, TCP/IP protocols) the same.

1.1.3 SAN topologies

Fibre Channel interconnects nodes using three physical topologies that can have variants. Topologies include:

- Point-to-point
- Loop
- Switched

The point-to-point topology consists of a single connection between two nodes. All the bandwidth is dedicated for these two nodes.

In the loop topology, the bandwidth is shared between all the nodes connected to the loop. Hub technology is generally introduced in this technology to allow for bringing nodes in and out of the loop easily. A major drawback in this technology is the Loop Initialization Process (LIP), which needs to occur every time a node is removed from the loop or a new node is added to the loop. The LIP process disrupts the loop and can cause data disruptions.

A switch allows multiple concurrent connections between nodes. This is generally referred to as a switched fabric. Switched fabric is more reliable than a loop because there is no need for the LIP process. New nodes can be added and removed without any disruption to the Fibre Channel processes.

1.1.3.1 Services

Fibre Channel services provide a set of functions, some of them required by the Fibre Channel protocol, and some of which provide optional enhancements to Fibre Channel's basic protocols. An example of a required protocol is the Login Server, used to discover operating characteristics of nodes. An example of an optional protocol is a Name Server that can translate IP addresses and symbolic names to port ids.

1.2 Why use a SAN for tape ?

Since most of the material written about SANs is about disk exploitation, we felt it necessary to explain how SANs fit into the tape environment.

1.2.1 Tape usage in Open Systems environments

To understand this, we really need to look into the history of tape usage in the Open Systems environment. Historically, in the UNIX environment, users in a multi-user UNIX machine could not allocate and write to a tape drive. Tape was reserved solely for the system administrator to create backup data.

Simple commands, such as the TAR command, were available to create backups of the system. Users were not allowed to take backups to tapes, because tape drives were expensive, and system software was not available to allow for the sharing of the tape drives among the various users of the system.

Recently, a new paradigm for backing up systems has been developed and marketed — this is called the network backup model. In this model (shown in Figure 6), used by most of the popular backup solutions (Tivoli Storage Manager, Legato Networker, Veritas Net Backup, EDM, Cheyenne ArcServe, etc.), the data to be backed up is passed over the LAN to a centralized backup server. This server then manages the tape drives and library, so that each client server need not have an attached tape drive.

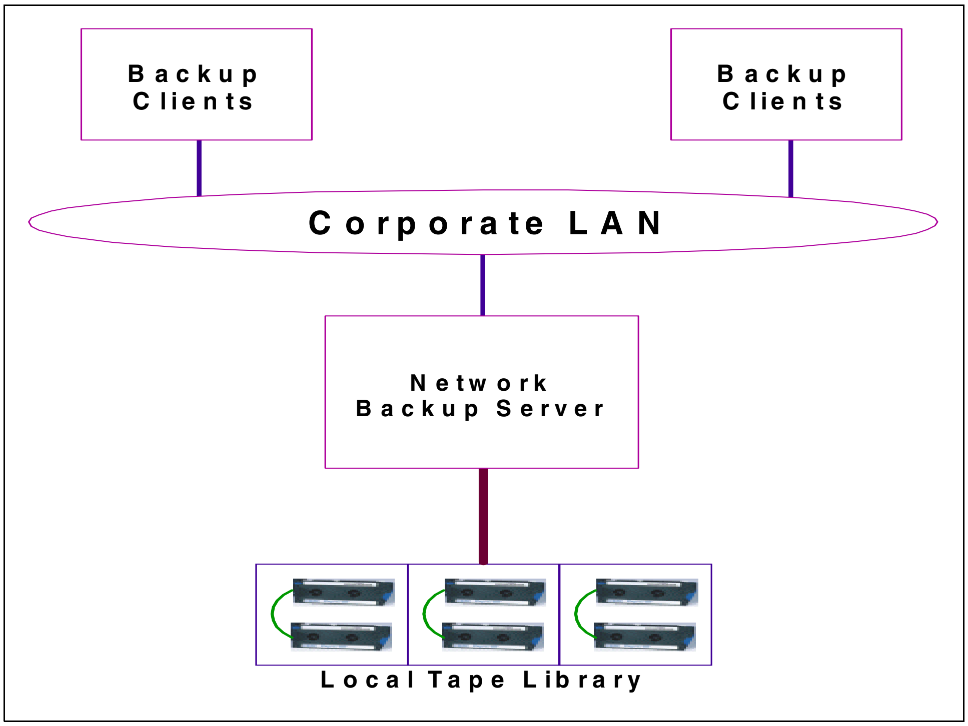

Figure 6. New network backup paradigm

This paradigm simplifies the backup process, makes the process more reliable, and generally lowers the overall cost of the backup process, since fewer tape drives are needed. Good, reliable tape drives (like the Magstar 3590) are expensive, and any solution that can reduce the number of these expensive drives draws the attention of the marketplace.

However, one type of client server did not fit neatly into this new paradigm — the large database server and other applications using large files. For this type of client, incremental backups are generally not practical, and full backups take a lot of network bandwidth, thus we often allow the large database server to retain access to their own tape drives, with their own copy of the backup package running on their server. This increases the requirements for tape drives, which increases the cost.

Now, the new technology, SAN, has positioned itself to solve this backup problem. Backup solutions can utilize the SAN technology in a number of ways to reduce the costs of their solutions at the same time they increase the performance of the solution. Sounds great — and it is! First, SANs allow the storage (tape drives) to be attached directly to multiple servers. Thus, if software can be developed to allow the sharing of these drives at a logical level, we will get the best of both worlds — direct connection of the tape drive to the server for high speed while sharing the drive with another server to save money. All we need is the proper traffic cop.

Interestingly, since almost all the tape mounts in the open world are requested by the backup software, we have the logical choice for the software solution. The advantage to the backup software is that it can turn a heterogeneous situation with multiple servers of different types into a homogeneous solution. The software simply makes it look homogeneous.

Therefore, if the backup software vendors wrote interfaces to allow tapes to be shared between these backup servers (if the access is allowed by the physical topology), we would have a wonderful situation. Well, they have started this process, and many of the vendors have either announced products to do this, or announced the intent to deliver such a product. IBM has made tape sharing available for Tivoli Storage Manager 3.7 and higher, with support available for Windows NT and SUN with Magstar tape drives in October 1999.

This is most notable in the NT environments, where IBM, Dell, and Compaq are rushing to beat each other to announce these tape SAN solutions. For example, the SAN Data Gateway from IBM allows servers to share a set of SCSI tape drives. These tape solutions meet and solve a real cost problem facing many customers today. The gateway today allows six hosts to attach and share four SCSI strings of tapes, and this can prove to be a really economical solution, if it allows the customer to reduce the number of 3590 tape drives or to share the same number of drives with more hosts.

1.2.2 Tape solutions are ready for the environment today

There is a fundamental difference between disk and tape which cannot be overstated — tape is a *removable* media. Thus a disk drive includes both the media and the read/write head; whereas, these two elements are separated in the tape environment. The world wants to share disk data and it knows that tape data cannot be shared (read by two systems at the same time). Therefore, the promise (and hope) for the disk storage world is to get to file level data sharing. This will takes years to achieve, if ever. Thus the hype on disk will fall short of the hopes of the customers during the next several years.

Sharing in the tape world comes in three distinct flavors:

1. Library sharing
2. Drive sharing
3. Media sharing

In many instances the tape library can be shared today. For the 3494 product this is part of its basic function. The library is shared in a physical way with each system thinking it really owns the entire library. The library manager ensures that integrity is enforced. Drive sharing is very difficult to accomplish without a SAN. In the S/390 world, where the ESCON Director can be considered a SAN, sharing is easy to implement. S/390 software products exist to ensure integrity of the drive. Sharing in the SCSI world needs the "any-to-any" features of the Fibre Channel technology and software. Media sharing is really not possible except to pass the tape between systems. However, sharing of the scratch pool can be very beneficial for resource utilization.

Nevertheless, since true tape media sharing is not expected, the very possibility of sharing tape drives will be a welcome relief for the customer. Tape drive sharing is possible today when the tape drives are connected to servers via a SAN. Since a reliable tape drive is an quite expensive device, sharing is going to be a very important economic factor. Today many enterprises opt for DLT drives due to the cost factor of dedicated drives. If several large systems could share one Magstar 3590 tape drive, the economics might twist to the 3590 tape drive's favor. Sharing of the tape drive will do for tape what enterprise storage did for the disk world.

Refer to 1.2.6, "Tape is expensive — disk is cheap" on page 13 for more details on the economics of tape processing.

1.2.3 Gigabit ethernet versus the SAN

The question still lingers about gigabit ethernet — Will it be fast enough to stave off the intrusion of FC and the SAN? Although this is a good question and makes sense to ask, for the major use of the tape drive in the Open Systems world, the question tends to be irrelevant. The paradigm of centralized backups to a backup server, with the data being served to the backup server over a gigabit ethernet, makes a great picture, but the TCP/IP overhead on the backup server is already too great. In a busy Tivoli Storage Manager server, the TCP/IP workload is generally about 50% of the entire workload. The SAN technology removes this large overhead by passing the data from one storage unit to another without the use of TCP/IP and its overhead. Therefore, even if the gigabit ethernet could pass the data as quickly, the reduction in processor overhead outweighs any perceived advantage of the gigabit ethernet solution.

Thus, not only can SANs be used today to solve the tape sharing problems, there is a real benefit to the backup server in reduction of processing power needed to do the backup work. This will allow each backup server to actually manage more clients.

1.2.3.1 What about NAS?

Network Attached Servers (NASs) were the rage several years ago. Connection of a large file server to the LAN network using NFS mounts is attractive for the file server environment. The client server simply goes to a very robust, special purpose file server for its files. These files are then delivered over the LAN. If the LAN were a gigabit ethernet, this would provide excellent service for the client. However, the use of these NAS devices has never ventured successfully into the tape world, and with the ascent of the SAN technology, it probably will not. We believe that the explanation given above, about TCP/IP, is the primary reason they will not be developed.

1.2.4 SCSI disk solutions can look like a SAN solution

The hype over disk SAN solutions may be overstated. Many storage vendors (IBM, Compaq, EMC, etc.) already market large scale SCSI based disk subsystems that offer many of the benefits that the SAN technology promises. The IBM ESS products allow multiple hosts to attach to the disk subsystem, offer some measure of shareability and connectibility, offer hardware mirroring features, and hence could be considered to be a rudimentary SAN. However, by our definition, since FC does not need to be involved in these multi-host subsystems, they do not constitute a SAN, but they can attach to a SAN. These vendors have been able to stretch SCSI to accommodate many of the requirements of the customers.

Thus, in many ways, the need to have the SAN involved in the disk subsystems might be overstated, since many of the problems can be solved with these SCSI subsystems. Of course, as SAN matures, the more robust FC architecture will allow the vendors to attach these SCSI subsystems to larger topologies and bring the promise of the SAN to this environment.

1.2.5 Tape solutions really need Fibre Channel

The major restriction for almost all SCSI tape drives (except the Magstar 3590) is the fact that they all have only one SCSI connection per drive. This limits the ability to create a second path to the device. Sharing of a drive would require multiple hosts on the SCSI bus, and this is difficult to manage and generally not recommended. Fibre Channel removes this limitation, since the very nature of FC is to have all nodes in the topology able to talk to all the other nodes. Thus even the smallest gateway can create savings in the tape environment and becomes a very marketable device.

1.2.6 Tape is expensive — disk is cheap

Have you ever heard this before? Probably not. In fact, you have probably heard just the opposite. But the fact is that writing to a tape or reading from a tape is very expensive, while storing the tape media is very inexpensive. Likewise, disk is very cheap to write to (and rewrite to) but expensive to store data for a long time. Figure 7 depicts how the costs change over time to store a gigabyte of data on disk and tape. Tape starts high, primarily due to the high tape drive cost, but the cost per day is very low; whereas, disk starts low and the cost per day is relatively high. Another factor making tape expensive is that often the tape drive cannot be shared with other hosts, and the drive can only be used for a small part of the day.

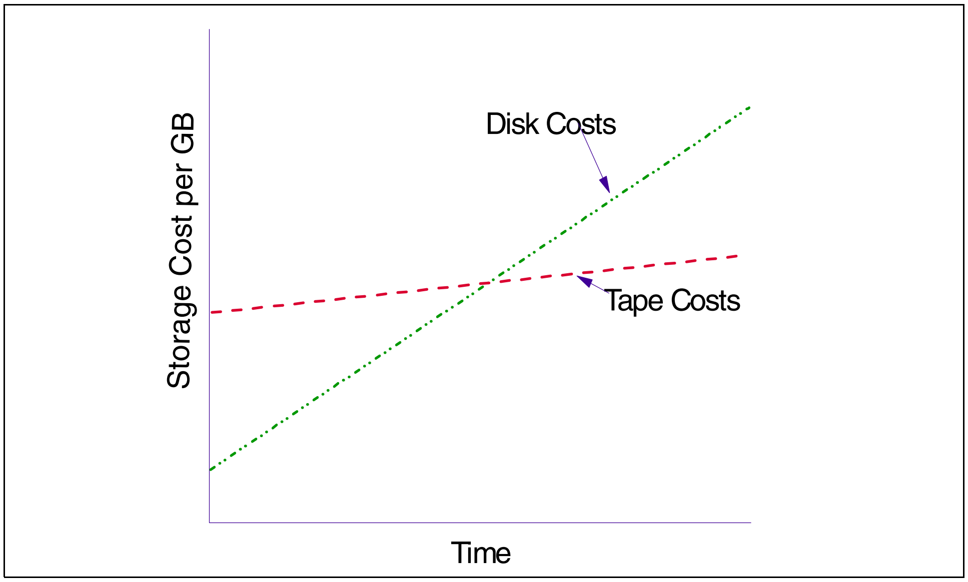

Figure 7. Disk vs. tape storage costs over time

The new backup paradigm alleviates this problem somewhat, especially for Tivoli Storage Manager with its intermediate disk storage (which is inexpensive for short term storage), but the database servers cause a breakdown of the economic picture. The SAN technology will allow these tape drives to be shared between backup servers while speeding up the process. Thus each tape drive can do more work; hence fewer are needed, and the cost per tape mount drops.

As you can see, if we can lower the cost of the tape mount, we can really effect the total cost of the tape backup, since much of the cost comes from the mount. Figure 8 depicts how the SAN technology lowers the cost of a mount by allowing tape sharing. The examples shown are for Tivoli Storage Manager, but they would be similar for any other network backup solution. Since the database servers own the tape drive, these drives will have fewer mounts per day, and hence have a higher cost per mount. The normal paradigm of Tivoli Storage Manager for fileserver type clients increases the mounts per day, thus lowering the cost per mount. Since the SAN technology allows us to share these drives across these two environments (while actually increasing the performance) one would expect even more mounts per day per drive, thus driving the cost per mount down even further. Note that Tivoli Storage Manager with its disk cache already increases the productivity of the tape drive by allowing data to move to tape over a longer period of time at a higher rate of speed.

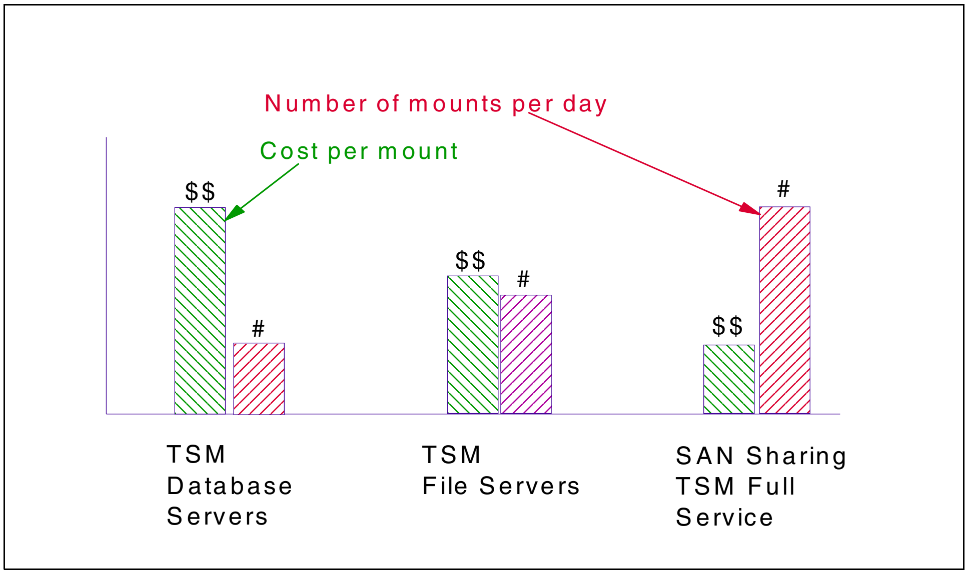

Figure 8. Cost of tape

1.3 SAN issues

Of course, with the new technology come some issues of implementation and usability. One major issue facing all the vendors is that of interoperability. Since the Fibre Channel standard encompasses a lot of technology (from disk use to tape use to video to printers), each with its own set of requirements, the early solutions put out by the various vendors did not always work when put with another vendor's product. As the industry matures and more testing is completed and fixes applied, we expect that there will be more and more devices that work together. This really was to be expected, and most vendors set up interoperability labs to test their products. IBM is operating three such labs, one for product authorization and two for customers to come to, for testing various cross vendor solutions with IBM's already authorized solutions.

Refer to the Web site of the IBM National Testing Center — Gaithersburg, Maryland at:

```
http://www.storage.ibm.com/ibmsan/sanlab.htm
```

One major issue that has not been resolved is how to use the various hubs and FC-AL in productive workloads. While testing the SAN Data Gateway, IBM discovered that an FC-AL hub has to go through LIP processing every time a node is either brought offline or online to the loop.This LIP processing often causes the loop to drop an in-transit I/O, which is not a good thing for the I/O process. In the case of hosts, this happens frequently. Thus IBM has decided that for the initial rollout of this SAN product, high-end platforms will not support the hub with multiple hosts attached to it. Since devices seem to be much more stable and disruptions of the loop are much less frequent with multiple devices attached, IBM will support the hub with devices attached to the gateway, and will support Windows NT hosts.

Sharing of the tape drives, which can now be seen by each of the attached hosts, can only be effective if software is available to manage the process of sharing between the servers. The paradigm that has emerged (used by several vendors) is to allow one of the servers to be the master and the other a slave. The slave asks the master for permission to use the drive (via the LAN) and when the master knows that the drive is free, permission is granted. Given homogeneous software platforms (all servers being either Tivoli Storage Manager, Legato, Veritas, etc.), this sharing can take place across heterogeneous hardware platforms.

1.4 Marketplace

Data is the most critical resource in today’s environment. Without data, computer processes would be very uninteresting. Data created today is used tomorrow and the next day. Thus the amount of data seems to grow without bounds in today’s enterprise. This has many causes:

- Dramatically lower costs of storage subsystems
- More complex data formats (see Figure 9), an evolution from:
  - Text
  - Spreadsheets
  - Images
  - Presentations
  - Video
- More complex applications, especially data warehousing applications

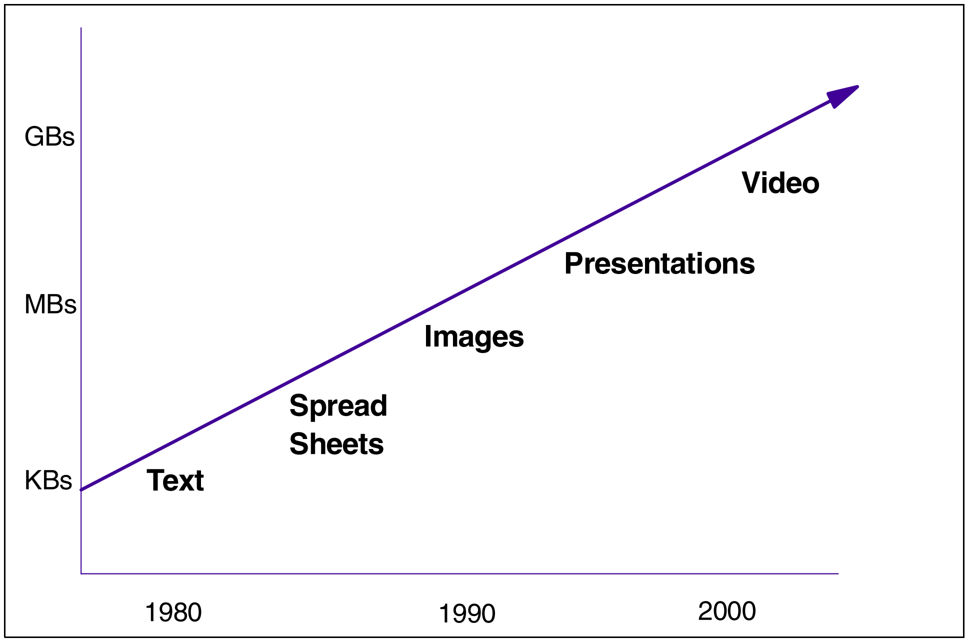

Figure 9. Exponential growth of data

Many enterprises are just starting to really understand how valuable their data is and use it to create new business and markets. Even at the PC level, this data is becoming critical, since it is the only place the data really resides. With the realization of the importance of the data, most enterprises come to realize that the current technology leaves them a bit short in terms of reliability, scalability, availability, accessibility, manageability and costs. As mentioned above, these issues have created the marketplace for the Fibre Channel architecture and the need for robust SAN topologies.

The major driver during the last few years has been to produce FC elements to solve the problems in the disk storage environments. It has only been recently that the vendors began to see, and then develop, solutions in the tape arena. We see the tape market being a very robust SAN market for the next several years, because the tape environment can be helped and improved so much by the inclusion of FC and SAN technology.

Chapter 2. SAN opportunities

As we can see from the discussions in the previous chapter, the advent of the FC technology and the SAN topologies will help resolve some fundamental issues facing those enterprises trying to stretch the limits of the older storage architectures. Hardware elements alone cannot solve all the issues; also needed are software management tools to really exploit these technologies. Vendors need to re-evaluate their existing products and possibly redesign the architecture to fundamentally use the benefits of the new technology to solve the problems their older products had tried to solve. By the very nature of an implemented product, it has to adhere to the available technology structure. When a new technology is available, restructuring is necessary in order to take advantage of the features of the technology. Simply patching the older software solution to work in the newer hardware technology will not yield maximum results for the product set. Those vendors who do this fundamental readjustment will be the short term winners.

So the question remains as to what kinds of products we can expect as the FC architecture and SAN topologies roll out during the next several years. In this chapter, we hope to show some generic solutions, including both those that are available today and those we might expect in the future. We will try to use generic terms such as UNIX server, NT server, S/390 server, tape library, backup server solution, etc. You, the reader, can put into the scenarios whatever product names you wish. Since this redbook is about the future of SAN technology and tape environments, we felt that this chapter is key to understanding why the vendors (and the industry) are pursuing these solutions so feverishly.

One should not expect that any one vendor can be the leader in every technical area. Rather, you should expect that the vendors will have to partner with others to deliver products to meet the needs of their customers. IBM’s SAN initiative is to utilize all resources at our disposal which will include close collaboration with partners as well as integration with other products from competitors. For example, expect that in the future IBM’s backup product Tivoli Storage Manager will deliver solutions on most other hardware platforms (just like today). So if Dell, for example, delivers a SAN solution which allows the sharing of tape drives, one might expect Tivoli Storage Manager to support the hardware solution as soon as possible. Another example is the FC concept of *third-party copy*, which enables *server-free backup*. Tivoli Storage Manager will exploit this feature as it becomes commonplace.

2.1 Solutions available in the near term

This section will discuss the solutions we see currently in the marketplace. We will keep this generic, but in most cases IBM has solutions that meet these needs. For the most part, this discussion is directed at solutions to problems, and not the elements used to implement those solutions. In later chapters, we will present the IBM solutions in detail with real configurations.

2.1.1 Tape connectivity

Connectivity to tape is essential for most backup processes. However, tape drives are expensive, and tape libraries are even more expensive. Study after study shows that automation of the tape process saves money while increasing reliability. Tape and disk are fundamentally different, since the media of the tape is removable, and storable anywhere, away from the tape drive. This is not the case for disk — the drive and media are locked together. Due to the removability and storability of tape media, enterprises have long had to utilize staff to remove these tapes, transport them to a storage site, and then return them to the tape drive for mounting when needed. Cutting out either expensive drives, libraries, or handling is generally the direction of customer tape planning initiatives.

The SAN allows for greater connectivity of the tape library and tape drives, especially at greater distances, and allows for eventual tape sharing. Tape sharing will be discussed in the next section. FC allows for 10 km between the server (or data point) and the connected tape node. This allows for both greater sharing and remote tape creation. Disaster recovery will be discussed below as well.

The single biggest issue today with tape sharing and remote creation is the SCSI architecture. Most SCSI tape drives currently have only one SCSI port and hence can only be on one SCSI bus. This severely limits the number of hosts that can physical allocate the drive without recabling. FC enables the multiple host scenario without recabling.

Figure 10 depicts multi-host tape sharing using a SAN. In this case three hosts are connected to a single SAN topology. Also attached to the SAN are six SCSI tape drives through a bridge or a gateway. If software to manage the sharing is not available, we would need to isolate (or zone) the drives to unique hosts. With the proper management software, each drive can talk to each host, and the connections can be dynamic without need for any recabling.

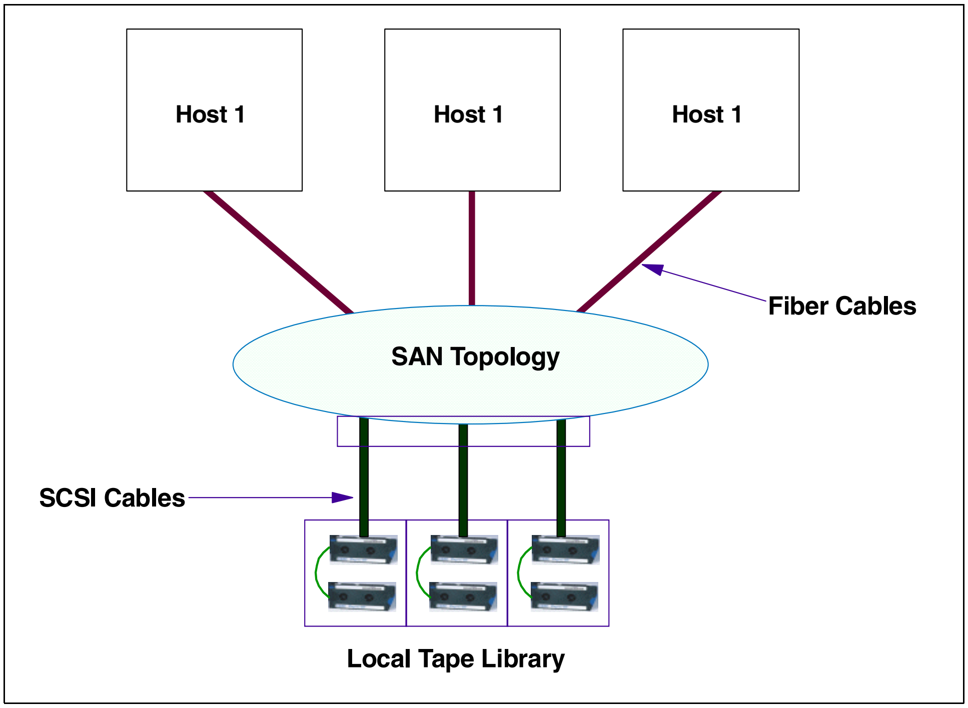

Figure 10. Multi-host tape drive sharing

2.1.2 Tape pooling or sharing

Tape sharing is one of the critical objectives of the implementation of a tape SAN, and also one of the complex ones: Tape sharing comes in different flavors and implementations. Tape resources can be split into three components, each of which can be shared:

- Tape library
- Tape drives
- Tape media

Library sharing has been practiced for some time using, for example, the multi-path feature on the IBM 3575, which allows you to partition a physical library into multiple logical libraries, or the library manager categories on the IBM 3494. (Refer to Table 2 on page 66 for more details on sharing and partitioning tape resources.)

The SAN allows multi-host connectivity to the drive, but software is necessary to allow for true sharing. Today most vendors are pursuing a model in which one server manages the tape resources and the other sharing servers need to ask permission to gain access to the tape resource. The real benefit here is the sharing of a very expensive resource — the large library and tape drives.

The rule of thumb for sharing resources is that for two systems sharing the same resource, there will be a reduction of about 20-30% in native resource requirements. For example, two Tivoli Storage Manager servers both needing four Magstar 3590 tape drives and 500 tape slots to manage their individual Tivoli Storage Manager workloads would be expected to need only 6 drives and about 800 slots to manage the same workload if they could share the resources. Thus, if the SAN technology elements and the new levels of software are not too expensive, as compared to what already exists, the savings in tape drive resources will generally provide a good economic justification.

Tape sharing solutions have been announced by several vendors, and this should become the launching point for many other more advanced solutions in the next few years.

Tape pooling is the ability to allow two or more servers to logically share tape drives within a tape library. In this case, the servers need to be attached to the same SAN as the tape drives, and there needs to be some software management to control who owns the drives and tape cartridges at any one point-in-time. Figure 11 shows three backup servers all getting their client data through the corporate LAN sharing six tape drives in one library. Software running in the master backup server talks with tape sharing client code on each slave backup server to control the allocations of the tape drives. The only difference between Figure 11 and Figure 10 is the software created by the backup server solution provider to actually manage the tape resources.

Tape pooling across a SAN will be supported by Tivoli Storage Manager, starting with release 3.7. The nature of the SAN is not critical to this software solution, since the control of the actual access is done through the Tivoli Storage Manager management software. Physical access to each tape drive is required, which SAN topology allows.

When using a tape library with outboard management, like the IBM 3494, the tape drives can be shared without one of the Tivoli Storage Manager servers taking the role of the arbiter: The tape drives can be shared using SCSI Reserve/Release functions (dubbed auto-share, a Magstar 3590 feature) and the tape media is controlled by the library and assigned to a specific host, thus eliminating the need to manage the tape inventory using application software.

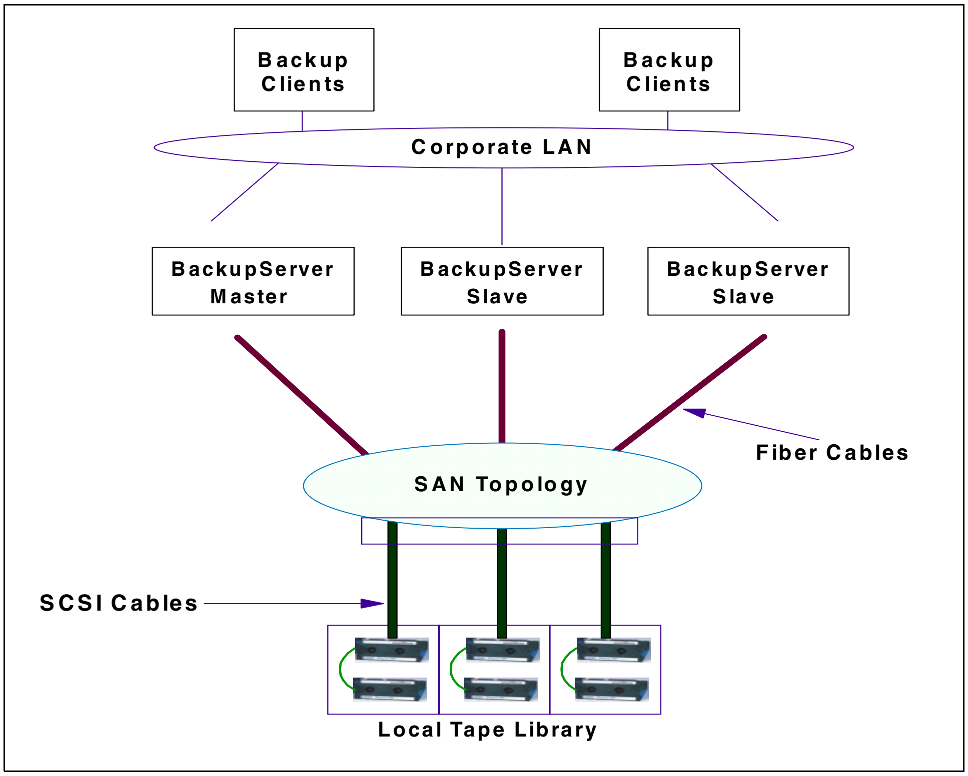

Figure 11. Tape pooling

2.1.3 Tape extension for disaster recovery

Another important aspect of tape SAN is the exploitation possibilities of tape connectivity for disaster recovery enhancements (Figure 12). Today most enterprises take their tape backups offsite for disaster recovery. The tape is actually created in a locally attached tape library, then ejected from the library and finally removed to an offsite location. All of this requires manual intervention and is error-prone. A major reason for failed recoveries is caused by the wrong tape being in the wrong place.

FC SANs allow the backup server to create the tape easily and safely in a remotely attached tape library. This removes all of the manual effort, since once the tape is created the tape is already offsite. Thus, for the disaster recovery scenario, this removes a major reason for disaster recovery failure.

A real need exists for extending tapes into a remote site. Currently, most enterprises need to staff the tape operations in order to handle the tapes as they come and go from offsite storage. Since Fibre Channel allows for greater distances, it becomes much easier to put a remote tape library across your campus to create the backup tape copy. With this method, no manual handling is necessary.

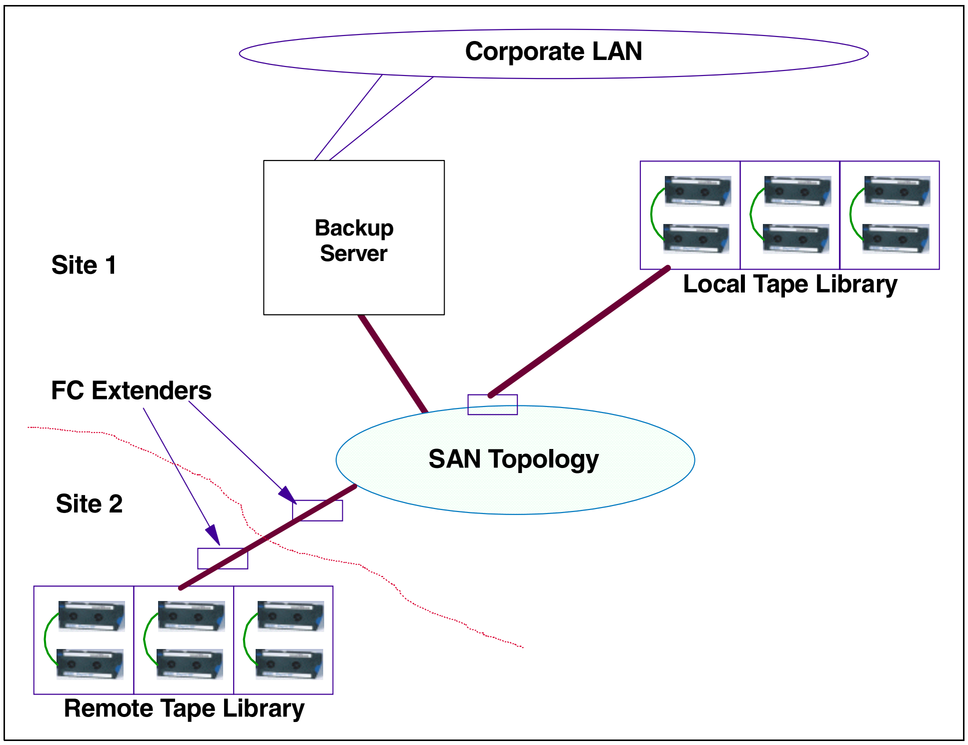

Figure 12. Remote tape creation using a SAN

2.1.4 Fully redundant, disaster tolerant tape solutions

As requirements for data protection increase, the value of these backups will also increase. Systems will be developed within an enterprise where the backup is a critical part of the success of the solution. Maybe instant point-in-time backups are required for legal reasons, maybe the tape information is actually needed for a real time information search; but whatever the reason, many enterprises will come to regard the availability of the tape drives and library as critical. A tape SAN can come to the rescue and can enable the necessary connectivity through multiple HBAs and multiple FC ports interconnected via the SAN. Elimination of every single point-of-failure can be assured via SAN topologies. For disk subsystems, this possibility has been available for years, and now it will be available for tapes.

In Figure 13, we show how a fully redundant tape subsystem might look using a SAN. In this case there are two SANs, each of which cross a geographic boundary, connecting two sites. If software is available to duplicate tapes automatically and to fail-over the servers, there would be no single point-of-failure. Local failover can be implemented where the tape subsystem and its components can be made to be fault tolerant and disaster tolerant.

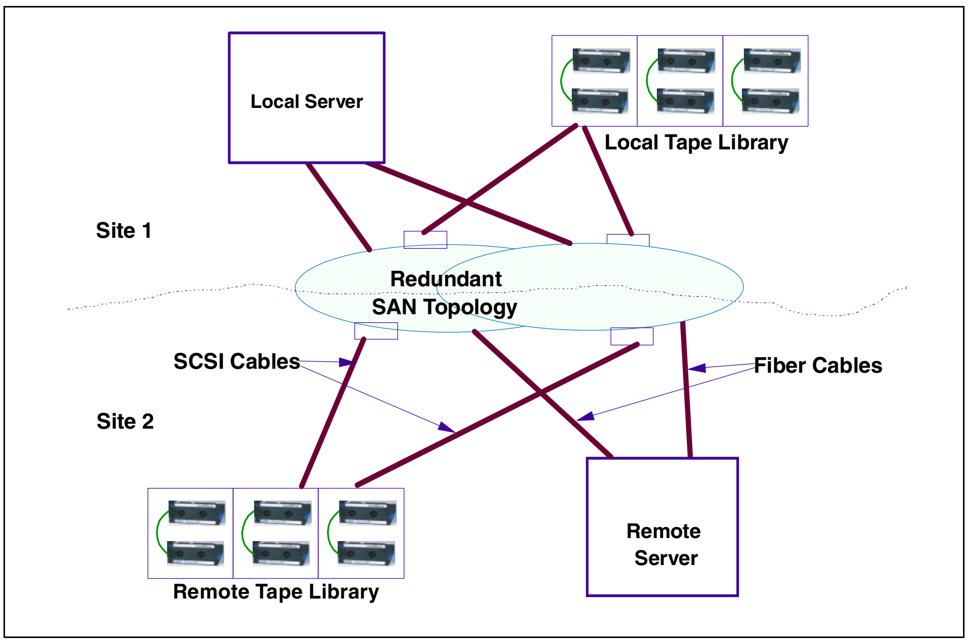

Figure 13. Fully redundant tape subsystem configuration

Although a fully redundant tape subsystem is not currently available, we place it in the near term solution portion because, once tape sharing and remote tape creation are available, a redundant subsystem can be built. Advanced features built upon this model will be discussed in the next chapter.

2.2 New opportunities

This section will discuss solutions we have only heard about or expect to see. Some of these are here simply to show that the new SAN topology with the underlining FC architecture allows us to think about new ways to solve the data management issues facing our customers today. It is these new potential solutions that are causing the keen interest in this arena.

2.2.1 LAN-free backup

Current LANs are too slow and often clogged. We need to remove the data flow from these networks so they can perform the interactive transactions that they were originally designed to do. Direct connections between the client and the server have been used in the past to remove the heaviest traffic from the LAN. Products like CLIOS and InfoSpeed allow a UNIX or NT client to connect directly to the S/390 server and move large blocks of data without using the corporate TCP/IP LAN network. These solutions were point-to-point utilizing the ESCON topology available to the S/390 server. This concept can be migrated into the general server world via the connection to a SAN topology. Without changing the software in the backup server, a SAN could be utilized to move data from the client to the backup server as long as the FC elements supported TCP/IP on the upper layers. However, the HBA vendors have not perfected this interface and the backup server software vendors are still waiting.

In addition to these issues, we have already discussed the heavy TCP/IP overhead within a backup server. No matter how the TCP/IP traffic arrives at the server, this overhead will exist. However, the FC architecture actually enables an even better solution where the data can flow directly from the client to the tape drive, which can be managed by the backup server. This is the concept of tape sharing with a software package, as discussed in Section 2.1.2, “Tape pooling or sharing” on page 21. There are actually many ways to implement this type of function but, in general, the tape SAN is crucial to the operation.

Tivoli Storage Manager supports a LAN-free backup mode of operation using Version 3.7 or higher. A LAN-free backup is done by a backup server using the SAN topology and functions of FC to move the backup data over the SAN, thus eliminating the LAN from the data flow. This does two things: first, the LAN traffic is reduced, and secondly (and most importantly) the traffic through the backup server is reduced. This traffic generally is processor-heavy because of TCP/IP translations. Here the backup server orchestrates the data movement, manages the tape library and drives, and tells the clients what data to move. The client does not have its data on the SAN, but it needs to be connected to the SAN. (See Figure 14.)

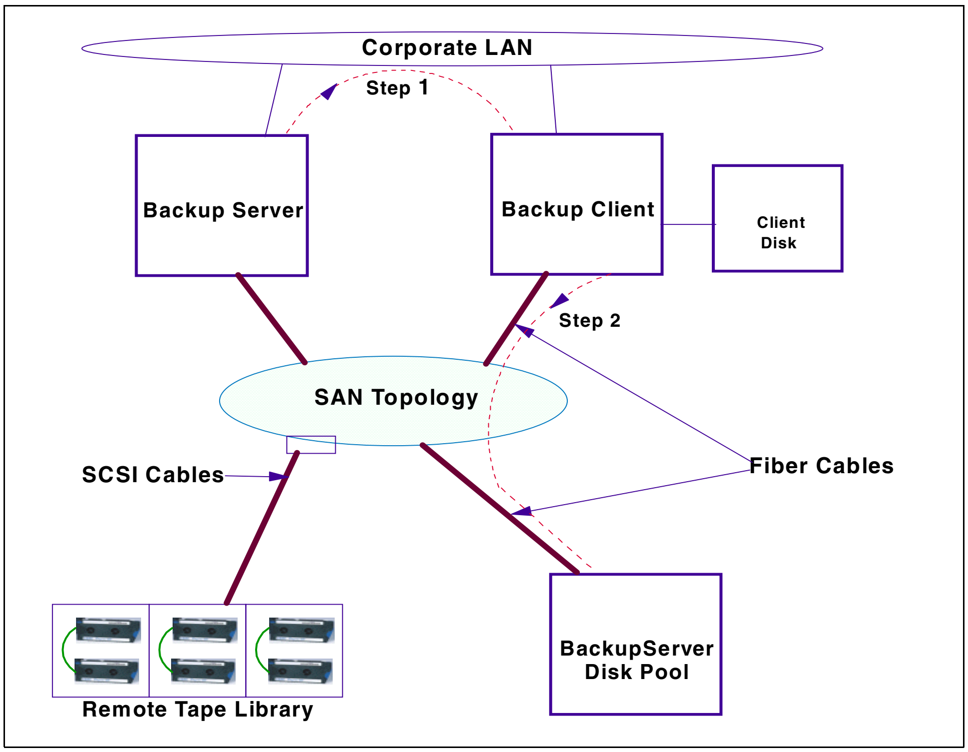

Figure 14. LAN-free backup

In Figure 14 we depict a LAN-free backup. Actually, a better term might be a LAN-lite backup, since the LAN is still used to pass meta data back and forth between the backup server and the client. However, the actual backup data is passed over the SAN. The meta data is the data needed by the backup server to manage the entire backup process, and includes things like the file name, the file location, time of the data movement, where the new copy resides, etc.

This data tends to be small as compared to the actual client data being moved.

In Step 1 the backup server talks with the client to determine what data needs to be moved; such as a backup, a restore or an archive. After the backup server determines what it wants done with the data, it sends the request to the client. In this case we show a request to move the data to the backup disk storage pool, such as might happen in a Tivoli Storage Manager environment. The client's own data is not part of the SAN topology in this case, so the client (in step 2) simply moves the data from its own disk to the backup server's storage pool as instructed to do. When the move is complete, the client notifies the backup server. This reduces a lot of system processor overhead on the backup server, since for a move through the server, about half of the server's processing power is used to process the TCP/IP transaction.

For this example we have assumed that the backup server uses its normal (non-SAN) processes of moving the data from the storage pool to the tape. Later, we will show how this is further improved when the backup server utilizes the SAN directly for its own workload.

2.2.2 Server-free backup

On June 22, 1999, IBM announced that Tivoli Storage Manager will support a server-free backup mode of operation in mid-2000. Server-free backup (see Figure 15) refers to the concept of using a feature which allows a storage device to talk to another storage device (such as a disk talking to tape directly). A backup server would need to understand what resources are available, would need to know where the data to be copied was, and where the data was supposed to be delivered. SCSI architecture allows for what is called third-party copy, where one SAN node delivers instructions to a second node to copy some data from that node to another node. The transaction takes place outside of the scope of the first node, and at completion, the primary node is notified of the success or failure of the data movement.

In this case, not only is the LAN not used for data transfer, the backup server is not even used. Precious I/O bandwidth and processor cycles are saved on the backup server and the client server. Many vendors are promising this functionality, and we expect to see this generally available from several vendors during the year 2000. This really will be a major evolution for the backup process.

Figure 15 is a rather complicated diagram showing the flow of data for an example of a server-free backup. This diagram is not intended to represent the way in which Tivoli Storage Manager will implement this feature, but rather to illustrate how one could implement the feature. In this case, neither the client nor the backup server is directly involved in the movement of the data. We have shown an example of a third-party copy where the meta-data for the actual move is passed to a data move manager (DMM), which is actually part of the SAN Topology (probably residing inside the switch or the gateway). The backup server works with the client to determine where the data should go, as we showed in Figure 14 on page 27, but now after the data is moved into the disk storage pool of the backup server, the power of the SAN is invoked.

It needs to be pointed out that the "data move manager" is a generic term coined by this redbook team and not intended to represent an actual product by any vendor. The term third-party copy is a standardized feature working its way through SNIA and requires the new SCSI-3 protocol. This SCSI-3 protocol is being developed expressly to be used in FC environments.

This example shows the client's disk still attached directly to the client. If the disk was attached to the client through this SAN, the backup server could actually initiate, through the DMM, a third-party copy of the data between the disk of the client and the storage pool of the backup server.

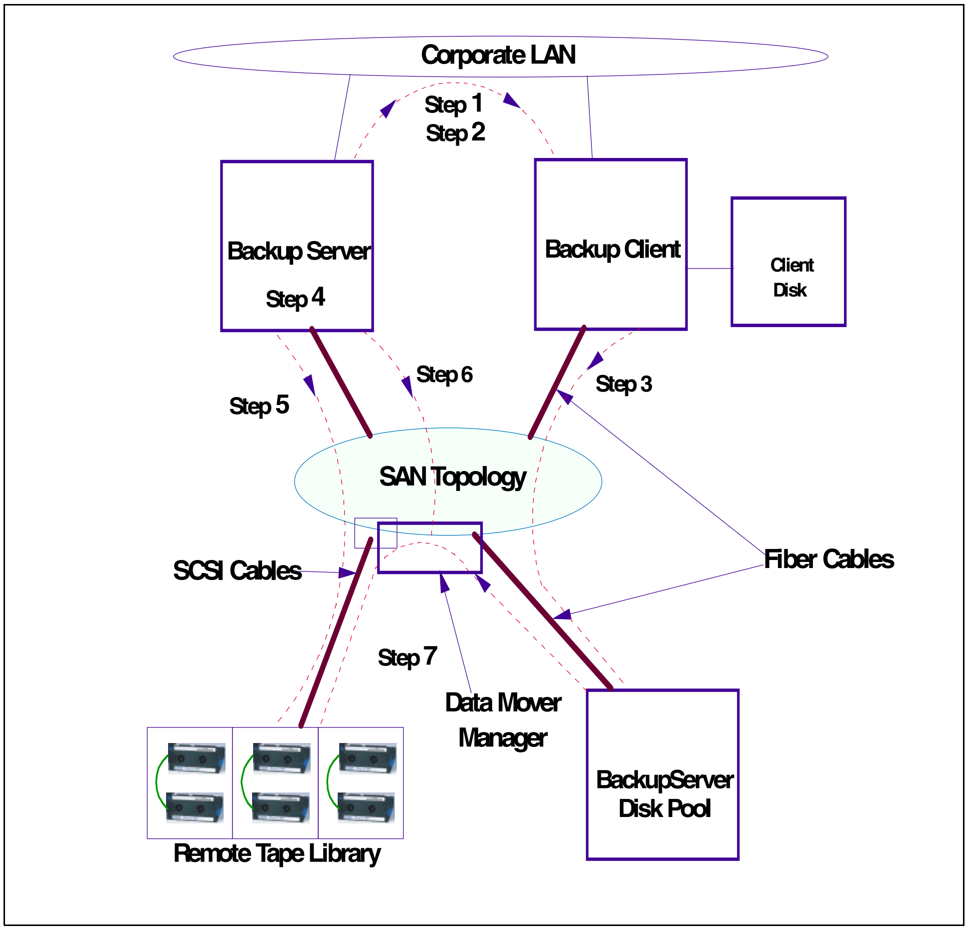

Figure 15. Server-free backup

- Step 1 — The backup server communicates with the server to determine what needs to be done; such as a recall, a backup, or an archive.
- Step 2 — The backup server communicates with the client to tell it to move the data (in this case) to the backup server's storage pool.
- Step 3 — The client moves the data from its own disk to the backup server's storage pool.
- Step 4 — At a later time, the backup server decides to move the data to tape.
- Step 5 — The backup server readies the tape drive, mounts the tape, and positions it.

- Step 6 — The backup server sends the third-party copy meta-data to the data move manager (DMM), which is located in the SAN topology.
- Step 7 — The movement of data is done without the involvement of the backup server by the DMM using the third-party copy feature.
- Step 8 — Notify the backup server of completion.

2.2.3 A new paradigm for a backup server using a SAN

In this section we will describe what the backup server of the future may look like, using an internal SAN. Creating such a server might actually be easier than creating servers that support LAN-free backup and/or server-free backup in a heterogeneous environment. For a single backup server, using a private SAN removes many of the issues of heterogeneity and establishes a homogeneous environment. We have assumed that the FC features of third-party copy and multicasting are used across this private SAN. Once these features are standard, we would expect the backup server vendors to create a server model like that shown in Figure 16.

We believe that the backup server vendors will need to exploit the SAN both internally and externally. For the next several years, large, complex SANs will be difficult to build, maintain, and manage. Therefore, we would expect that external SAN connections to the backup server will be reserved for selected large scale clients. Smaller clients will continue to pass their backup data over the LAN to the backup server. Thus the need for the backup server, especially a Tivoli Storage Manager backup server, to move data between disk and tape, and tape and tape, will remain. Removing this internal data movement from the Tivoli Storage Manager server will be of great importance to reduce the load on the server and increase its scalability.

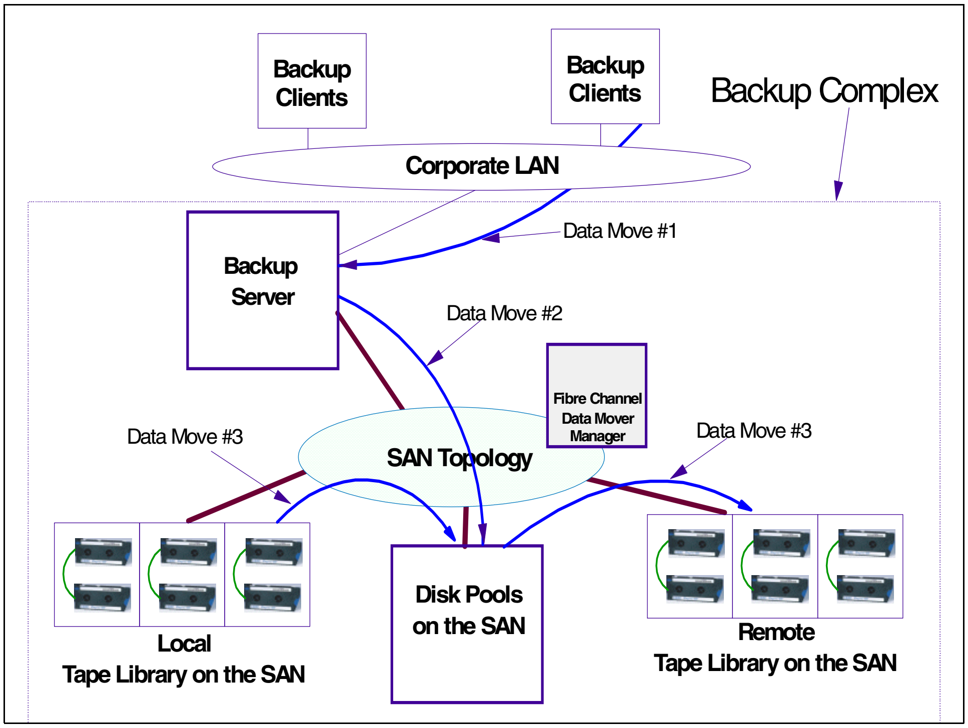

Figure 16. A new paradigm for the backup server using an internal SAN

In Figure 16, there are only three data movement steps even though there are four data movements. Steps 1 and 2 are managed by the backup server and the data flows through this server. Thus this is not classified as a LAN-free or server-free backup model. For smaller clients this will probably be the paradigm for a long time to come. However, by utilizing the concept of the DMM, Step 3 is managed by the SAN and reduces the workload on the backup server. Step 3 is classified as a third-party copy with multicasting. The backup server tells the DDM what to do and the DMM does the actual data movement. In this case making a remote copy of the tape data at the same time as the local copy. Thus two data movements are done in this one step.

2.2.4 Complex backup scenarios

The FC architecture allows for a SCSI-3 protocol feature called *multicasting*, where one FC command can direct the same output to two separate and distinct FC nodes. Thus, for the disaster recovery scenario, when coupled with the server-free backup depicted above in Figure 16, the tape local and backup copy can be made via one set of commands which uses no LAN bandwidth, backup server, or client resource, but rather utilizes an element of the SAN and the bandwidth of the SAN.

In this case, the backup is managed via a backup server. Data movement is done by a function of an element in the SAN. Backups are taken as they are today, but over a high speed transport, and tape copies are made both remotely and locally to satisfy disaster recovery requirements. See Figure 17.

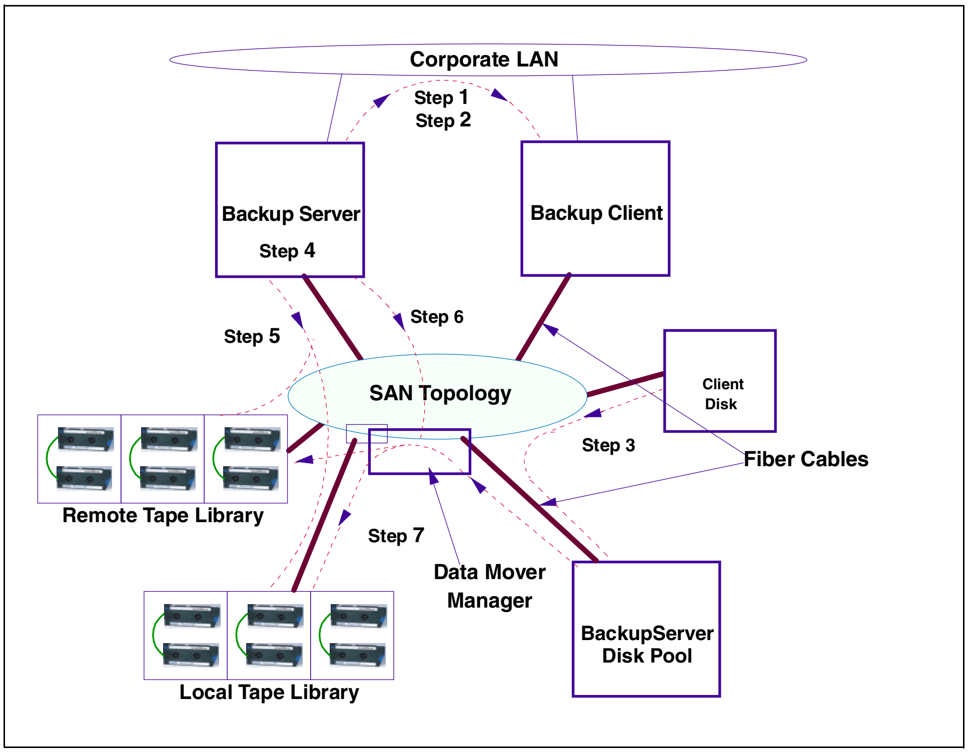

Figure 17. Server-free backup with client's disk on SAN and multicasting

- Step 1 — The backup server communicates with the server to determine what needs to be done; such as a recall, a backup or an archive.
- Step 2 — The backup server communicates with the client to tell it to move the data (in this case) to the backup server's storage pool.
- Step 3 — The client (or the backup server) moves the data from the client's disk to the backup server's storage pool using third-party copy.
- Step 4 — At a later time, the backup server decides to move the data to both a remote tape and a local tape.
- Step 5 — The backup server readies both tape drives, mounts the tapes and positions them.
- Step 6 — The backup server sends the third-party copy meta-data to the data move manager (DMM) with the multicasting feature enabled so as to invoke the dual movement of the data to both tape drives.
- Step 7 — The movement of data is done without the involvement of the backup server by the DMM using third-party copy feature of FC. Both tapes have the same data and look alike at this point.

Figure 17 depicts a very advanced backup server and client, both of which can take advantage of the SAN topology and FC architecture. Here the backup server works with the client to determine which data to move. Then, instead of the client actually doing the move, either the client or the backup server could send a request to the DMM, which would do the movement of the data from the client to the storage pool of the backup server. At a later time, the backup server then uses third-party copy and multicasting to move the data from the storage pool to two distinct tapes in two libraries at two sites. This is very advanced, because with this one set of commands, both the primary and backup copy are made to two tapes, and one is already off-site.

As an initial starting point, we would expect that the backup server will use the SAN technology first to deal with its own internal work. Due to the current size restrictions of switches and gateways, we would not expect a complete server-free solution for a while. However, the FC concept can be utilized by the backup server, and especially for a disk based system, like Tivoli Storage Manager, the movement of data along the SAN topology using third-party copy and multicasting can reap great benefits.

Figure 14 on page 27, Figure 15 on page 30, Figure 16 on page 32 and Figure 17 on page 33 have all depicted the backup server using the SAN functions of third-party copy. This relieves the burden of moving data from the server. All of these examples are things we might expect to see over the next few years. It is unlikely that these features will be available in the short term, since the rough edges of FC architecture and SANs still need to be smoothed over. Customers may want to begin planning with their backup vendor what the future will look like for them. Some vendors will lead and some will follow — IBM has made it clear that Tivoli Storage Manager will be leading the SAN charge. Over the next few years the vendors will begin to push out solutions that will exploit the SAN technology much like we have shown in these diagrams. You, as a customer, will probably need to begin to understand the implications to your environments, because these types of solutions are coming.

2.2.5 A new architecture for the backup server

The SAN attachment of the tape, the backup server, the larger clients, and the disk pool will allow (actually force) vendors to take another look at their products. We might see that the data movement is separated from the backup server inventory database. Today the backup server does three things: it manages the backup process and understands where the data is, manages the tape subsystem and media and understands what to mount, and moves the data. The SAN enables these three functions to be de-coupled and even moved into different servers. One server would need to be the brains and understand where things are and understand how to get them. A second server could manage the entire tape subsystem, the Tivoli Removable Media Manager server. A third server could manage the actual data movements, probably using FC third-party copy functions, as well as standard data movement functions. See Figure 18.

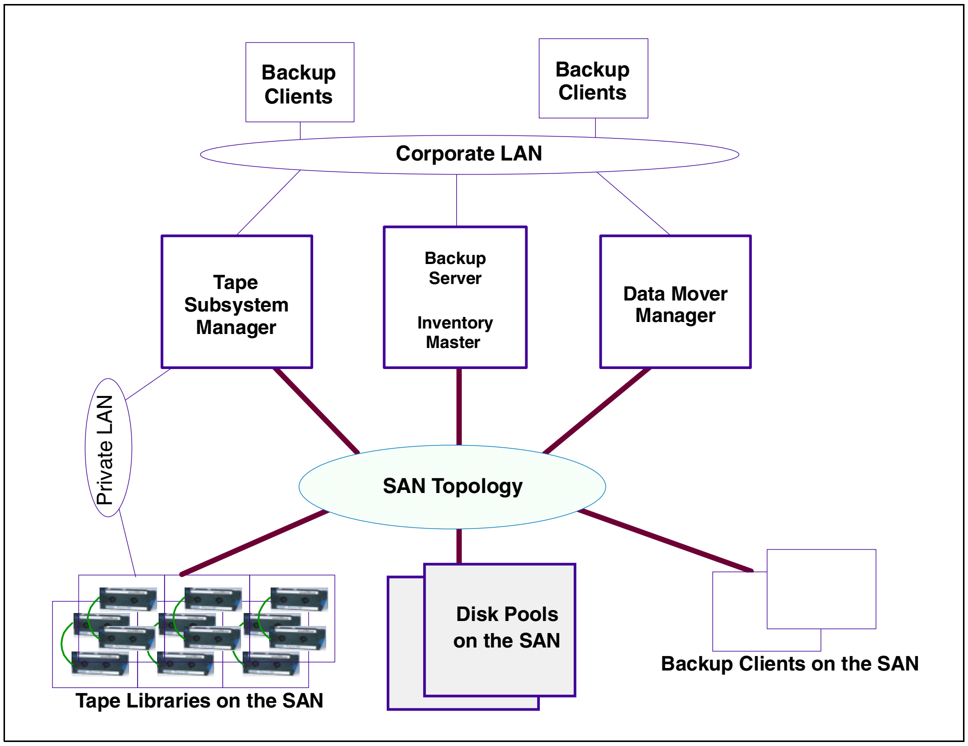

Figure 18. Backup architecture of the future using a SAN

2.2.6 Intelligent tape subsystems (ITS) and SAN — is this the future ?

The IBM Magstar Virtual Tape Server (VTS) has been enhanced to support attachment to SCSI host systems, currently limited to AIX. In addition, direct FC connection to the VTS will be allowed in the future, with FC attachment using the IBM SAN Data Gateway to be made available first.

The question has been asked repeatedly: why use a VTS for Open Systems connection? Clearly the issue of stacking small tapes onto larger tape media is more likely to be an S/390 issue, and clearly this was the major intent for the original VTS. But the current VTS is much more robust and has much more performance, and therefore can handle much bigger workloads. Virtualization saves real resources and, if the performance can be made good enough, can be used to reduce real costs to the enterprise.

Let us examine the VTS in much more detail to see how it could help solve some of today’s tape issues, and hence why it might be a good candidate to connect to a SAN topology. VTS is basically a system that appears to the user like many tape drives with many tape media volumes, but in reality it is a very large disk cache (to accept the data quickly from its clients), with enough real drives to off-load the cache when necessary and enough tape cartridges to store the logical volumes. If we recall some of the issues of tape versus disk (that tape is expensive to write to, but cheap to retain data on) and we map that to the way a VTS works, we see that if the disk cache was large enough to handle most of the tape recalls and expirations, then we would be putting *tape data* on disk for a short while (6-10 days) and then to tape for a very long while (over 30 days).

This method really maximizes the effectiveness of the two types of storage, without the client knowing, or caring, what is being done. Data kept for less than 10 days, which tends to include most of the active data and the databases, would really be going to disk, and not to tape, while data needing a long life would pass through the relatively inexpensive disk on its way to the real tape media. SInce most backup data is recalled in the first 10-14 days, if this data could somehow be in cache, no tape mounts would be required. So, properly configured, a VTS can be used to reduce the expense of tape mounts while giving a performance boost.

Since the VTS is based on IBM storage, the RS/6000 AIX processor, and Tivoli Storage Manager software, we should anticipate that a natural progression for the VTS is to incorporate the SAN technology under the covers. What does this do? It allows the VTS to continue to become even more robust, with even faster working pieces and larger storage caches and faster tape drives. As this becomes more expensive, we would really want to be able to share it. That is where the SAN comes in. Therefore, if the VTS continues to grow in robustness, size, performance, and features, the need for it to connect to a SAN will also intensify.

Can backup software solutions really use the VTS? Yes, in fact, the VTS is used in many Tivoli Storage Manager installations today. It is not necessarily the best fit today, but for individual enterprises the relatively low cost of sharing the Tivoli Storage Manager workload within a VTS with the other S/390 workload overshadows the cost of added native 3590 tapes which might make a better technical fit. For the other backup solutions, which do not use a disk cache, the use of a robust VTS could be a benefit. These other solutions work much better when the number of drives is increased — and this is precisely what the VTS creates. These solutions might actually work better with a VTS with 64 logical virtual drives (backed up by 1TB of disk and six 3590s) than with 16 or 32 real DLT tape drives.

Note: In order to differentiate the current IBM VTS product from possible future implementations, in this redbook we will call the VTS of the future the *Intelligent Tape Subsystem*, or *ITS*.

Can we develop other solutions, given plenty of tape and media access, quick recall times, and low overall costs? Would vendors discover new applications fashioned around an intelligent tape subsystem connected to lots of servers via the SAN? Could information retrieval systems currently utilizing optical storage devices be better served by the SAN-based ITS?

In Figure 19 on page 39, we have depicted what we might see in the future when vendors couple the concepts of intelligent tape subsystems with SAN technology. (Note that the current VTS marketed by IBM should not be confused with what we are showing here. IBM has not announced anything like this, nor would we expect to see this level of solution for several years.) In this case, intelligent tape subsystems are actually managed by an intelligent tape subsystem manager. Note that intelligent tape subsystems utilize a SAN to better use the real hardware. In this picture we have shown the subsystems sharing one real library at a remote site. Functions like third-party copy and multicasting would certainly be employed.

We would expect that new features may be developed that will enhance this picture even further. The net result is that the attached S/390 or Open Systems hosts (Host 1,2, and 3) believe that they are talking to real tapes, but they are really only talking to the intelligent tape subsystem. It is only the ITS that understands the physical relationship of the drives and tape media to the logical relationship presented to the tape client. This is analogous to the intelligent disk subsystems of today. Thus different physical environments can be maintained which might include a dual copy of each tape at a remote site without the real tape clients needing to know what is happening.

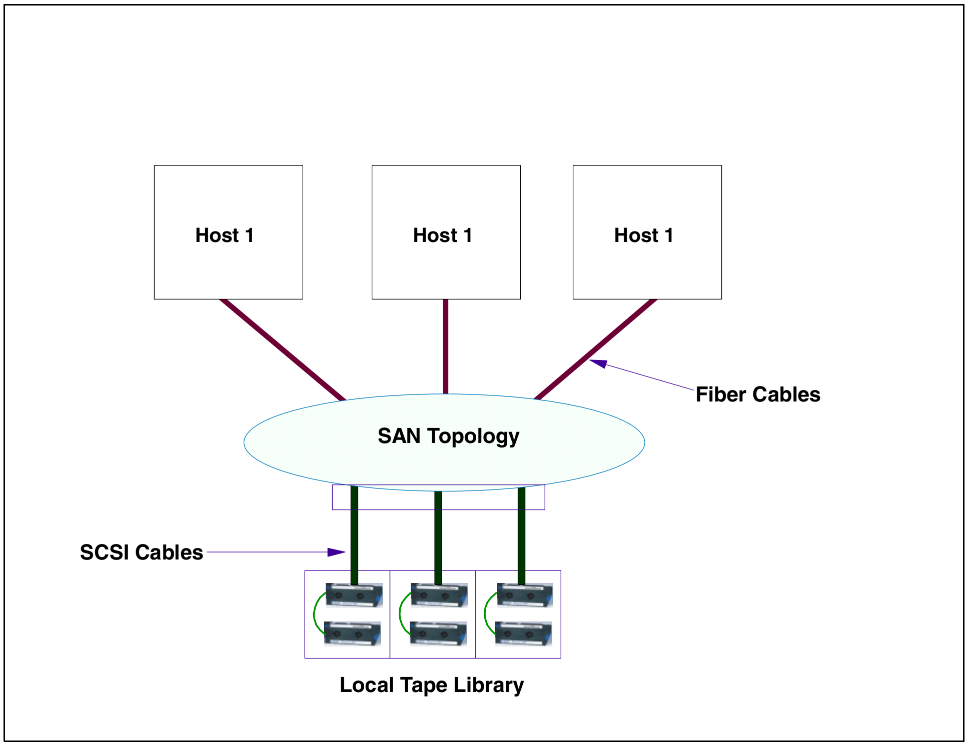

Figure 19. A SAN based virtualized tape subsystem for the future

As these robust ITS are developed, the tape processing world will never be the same again. The logic of moving data to tape or disk will become unnecessary, since the ITS creates the best of both worlds. Data is moved to the ITS, where it will actually reside initially on disk. After a while. it migrates to tape. However, if the data has a short life it may never make it to the tape environment. Thus cheap disk is used for short term storage and cheap tape is used for long term storage. Clearly it is the coupling of these two technologies (SAN and ITS) that allows this explosion in exploitation. Again we would caution the reader that these types of solutions are not yet around the corner. Work with IBM to understand the power of these technologies and begin to position yourself for this future.

Chapter 3. SAN solutions

This chapter is a review of the IBM products and intentions in the SAN area. We begin with a general overview of the IBM SAN Initiative statements, which summarize the IBM vision for SAN. Then, we provide a section for each individual area of the SAN Initiative, in which we describe each product or intention of IBM for that area.

3.1 IBM SAN Initiative

In this section, we give you an overview of how IBM intends to provide solutions and products in the SAN world.

IBM formally announced its SAN Initiative on February 17, 1999, and June 22, 1999. See Press Releases at:

```
http://www.storage.ibm.com/ibmsan
```

These announcements outline IBM’s vision and intent on providing solutions for the SAN environments:

> “Connectivity, management, exploitation and services make up the foundation of IBM SAN deployment. SANs combine connectivity hardware, such as Fibre Channel hubs, switches, and gateways, with software management that accommodates both IBM and non-IBM products and software solutions exploiting the storage, access, flow, and protection of information seamlessly, any time, anywhere. Leveraging its extensive IT planning, design, and implementation experience, IBM will provide the support, services, and education required to support end-to-end SAN solutions.”

Figure 20 outlines five key promises of IBM’s SAN Initiative:

1. Hardware base (servers and storage)
2. Connectivity
3. Management
4. Exploitation
5. Support

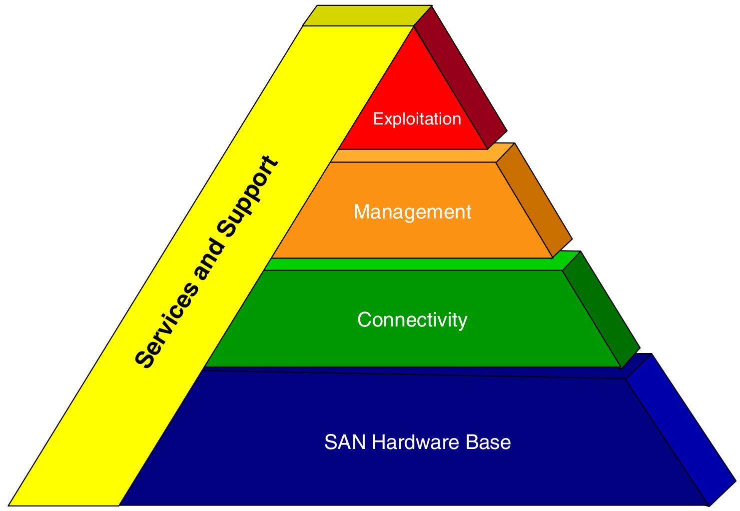

Figure 20. IBM SAN Initiative

Hardware base: IBM will introduce new storage units, and upgrade existing storage products, to allow you to integrate them in a SAN, either by providing native Fibre Channel attachment, or via gateways.

Connectivity: IBM will complement its current products with SAN-enabled products, from native FC-attached disk and tape storage devices to fabric components like gateways, hubs, and switches. In addition, migration paths will be provided to allow the use of existing devices.

Management: IBM will provide the management functions needed to configure, control, and monitor, not only the SAN elements (switches, gateways, etc.), but also the storage devices and hosts attached to the SAN. This will initially include management tools for individual devices, and ultimately a single interface to manage all the SAN resources.

Exploitation: IBM will provide applications and functions that will allow you to take advantage of the potential of SAN, including tape pooling, disk pooling, third-party copy services, file pooling, clustering, and data sharing.

Services: IBM support structure and services will provide the necessary support for the implementation of SANs in the enterprise. This includes education, consulting, installation, and maintenance services, as well as interoperability testing facilities.

These functions will be introduced in a staged manner. Some will be implemented using non-IBM developed products, notably for the SAN fabric components. For instance, the IBM SAN Data Gateway is a product from Pathlight Technology, and the IBM SAN Fibre Channel Switch is a product from Brocade Communications Systems.

See product information at:

```
http://www.storage.ibm.com/hardsoft/products/sangateway/sangateway.htm
http://www.storage.ibm.com/hardsoft/products/fcswitch/fcswitch.htm
```

Also, transition products (gateway products allowing to attach SCSI or SSA devices to a SAN) will be provided, to allow the integration of existing storage devices in SAN.

IBM is in a unique position to provide comprehensive SAN solutions, as there are not many companies on the market who can address all the elements of SAN, including support and services.

The IBM offerings will be open, and you will be able to integrate non-IBM hosts, devices, and applications, which is one of the main goals of SANs.

3.2 IBM SAN Hardware Base layer

In this section, we describe the products that are part of the SAN Hardware Base layer (Figure 21) of the IBM SAN Initiative. Here we talk about storage devices that can be attached, one way or another, to a SAN, concentrating on tape drives, of course, as they are the subject of this book.

The equipment on the other side of the SAN, namely processors, are also part of this Hardware Base layer. We are not describing host systems in this chapter, except to mention that these can be UNIX or NT machines, using the FCP protocol (SCSI), or mainframe S/390 systems talking FICON over Fibre Channel.

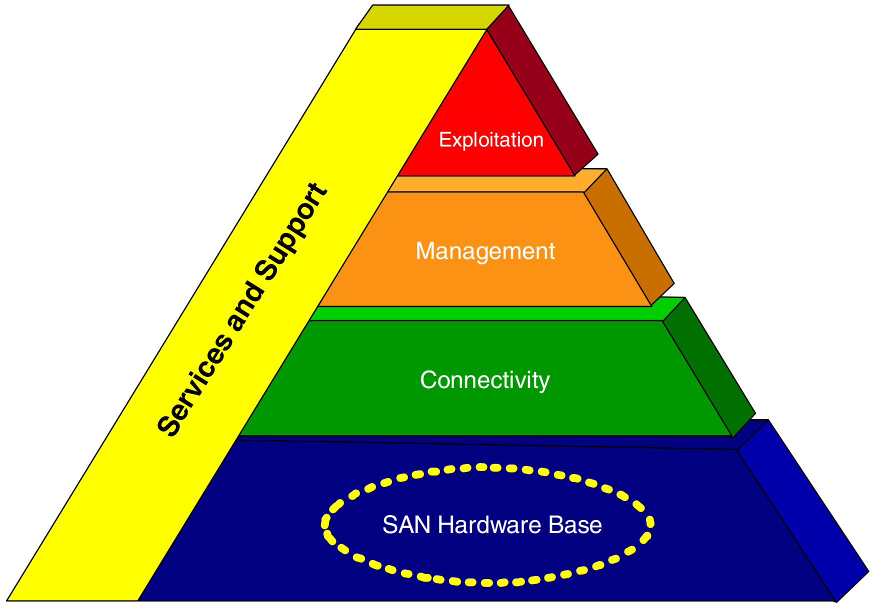

Figure 21. SAN Hardware Base layer of IBM SAN Initiative

3.2.1 IBM Magstar MP Product Line

The IBM Magstar MP 3570 Tape Subsystem is a midrange tape storage solution line of products that has the peculiarity of offering very fast access times to tape data.

This is possible in part because of the patented Magstar MP Fast Access Linear Tape Cartridge. The cartridge contains a self-enclosed tape path that eliminates time-consuming and cumbersome tape threading, loading, and unloading. In addition, the tape features an innovative new midpoint load design that dramatically reduces random search times.

The way this is achieved is by using a two-reel cartridge type. The Magstar MP 3570 attaches to AS/400, RS/6000, HP, Windows NT, Sun, and other SCSI-attached systems.

The Magstar MP 3570 can be integrated in a small library with a capacity of 20 cartridges slots and two drives, or 140 GB of data (420 GB with a 3:1 data compression ratio), or in a medium size tape library, the IBM Magstar 3575, which comes in several models, the largest one having a capacity of 320 cartridges slots, equivalent to 6.7 TB with a 3:1 compression ratio.

The Magstar MP drives delivers sustained data transfer performance of 7 MB/sec (native) and up to 15 MB/sec per drive or 100 GB/hour per two-drive library (with maximum compression).

The Magstar MP can be integrated in a SAN by connecting it to the IBM SAN Data Gateway. In addition to the increased distance and better cabling, this allows you to connect the 3570 to more than one host. This is something that you cannot do with a normal SCSI connection, unless you put more than one host on the bus, which is not recommended.

The number of hosts that can be connected to the 3570 is now a function of the SAN Data Gateway capabilities; currently, the limit is four hosts (the gateway has a maximum of four Fibre Channel ports). When the 3570 is attached to a router or a switch, you will be able to connect more hosts.

The Magstar MP 3575 library, if it has more than two drives, and depending on the library model, can be partitioned logically in up to 3 libraries, each one having a fixed number of cartridges slots. In a SAN environment, this functionality does not change. Through its inherent partitioning function, the 3575 library is especially suited for SAN environments, since SAN offers further connectivity and distance possibilities. So you can statically share this library between up to three hosts, which can be on different locations in your environment. Additionally, you can connect the library to even more hosts in a SAN environment and switch the logical libraries simply by reconfiguration of the SAN components.

3.2.2 IBM Magstar 3590 Product Line

The IBM Magstar 3590 tape subsystem is a family of tape drives designed to provide high capacity, high performance, high reliability, and connectivity for a wide range of hosts. The Magstar 3590 models B11 and B1A read and write in 128-track format. While reading or writing 16 tracks at a time, the Magstar 3590 uses serpentine, interleaved, longitudinal recording technology to make a total of four round trips from physical beginning to physical end of the tape and back again. The tape read/write head indexes, or moves vertically, when it completes each round trip so that the recorded tracks are interleaved across the width of the tape.

The Magstar 3590 uses a metal particle medium in the tape cartridge that can store 10 GB of uncompacted data. The integrated control unit uses a compaction algorithm that can increase the storage capacity of a single cartridge to 30 GB.

The Magstar 3590 Models E11 and E1A provide increased capacity and improved performance over the B11 and B1A. The E models have a native data transfer rate of 14 MB/sec, which is an increase of over 50%. In addition, the new models write in 256 track mode and double the capacity of the standard 3590 cartridge to 20 GB (uncompressed). Finally, the Magstar 3590 cartridges are read compatible with the E drive models. This means that any cartridge written by the B models can be read by the new E models.

IBM is providing investment protection to existing Magstar 3590 customers by making the existing B drive models upgradeable to the new E models. In addition, the B models have been repriced to make them exceptionally attractive to customers who may not need the added capacity and performance of the new E models.

The Magstar 3590 tape drive can be connected to a SAN using the SAN Data Gateway. The 3590 offers a particularity not found on most of the other tape drives: each drive has two SCSI ports, allowing it to be connected to two different hosts. The twin-tail feature of the 3590 extends the number of hosts that can be connected to a 3590, as it can be reached through two SAN Data Gateway boxes, and the unique 3590 auto-share feature allows reservation of the drive on the device level.

3.2.3 Fibre Channel RAID Storage Server

The IBM Fibre Channel RAID Storage Server is a robust storage solution for environments that require Fibre Channel attachment. This storage system can be attached to servers running AIX*, Sun Solaris, HP/UX, Windows NT**, or Novell NetWare — and it can be shared simultaneously by two operating systems, such as UNIX** and Windows NT. Multiple server attachments can be implemented via the attachment of the IBM SAN Fibre Channel Switch or the IBM Fibre Channel Storage Hub.

For enterprises with multiple platforms that share a storage system, this configuration can help protect storage and server investments while potentially reducing the overall cost of ownership. The Fibre Channel RAID Storage Server supports heterogeneous environments by enabling each RAID controller to work independently.

3.2.3.1 High-availability design

The Fibre Channel RAID Storage Server has dual-active RAID controllers that provide high throughput and redundancy. Both controllers can be simultaneously active to provide seamless failover capability in case of emergency.

To increase availability, each RAID controller supports up to 256 MB of battery-backed cache and can mirror write operations. Dual fans and power supplies further support 24x7 operations.

3.2.3.2 Extended distances

The IBM SAN Fibre Channel Switch and the IBM Fibre Channel Storage Hub provide greater flexibility and extended distances across Fibre Channel SAN topologies — enabling the Fibre Channel RAID Storage Server to be configured at distances of up to 500 meters via short-wave fiber connections or up to 10 km via long-wave fiber connections.

3.2.3.3 High performance

The Fibre Channel RAID Storage Server has dual Fibre Channel ports that provide an aggregate bandwidth of 200 MB/sec. With dual Fibre Channel Windows NT hosts, sustained data transfer rates of up to 179 MB/sec can be achieved with 256 KB sequential reads (up to 190 MB/sec for 256 KB sequential writes for 100 percent cache hits). Approximately 16,200 I/O operations per second may be reached with 4 KB 100 percent cache hits in certain environments.

3.2.3.4 Dynamic management

The StorWatch Fibre Channel RAID Specialist is a network-based integrated storage management tool that helps storage administrators configure, monitor, dynamically change, and manage multiple Fibre Channel RAID Storage Servers from a single Windows** 95 or Windows NT workstation. High availability and full redundancy are provided with the host-specific Fibre Channel Storage Manager software, which resides on the host system and provides automatic I/O path failover when a host adapter, IBM SAN Fibre Channel Switch, IBM Fibre Channel Storage Hub, or a storage controller fails. See product information at:

`http://www.storage.ibm.com/hardsoft/products/fcss/fcss.htm`

3.2.4 Enterprise Storage Server

The IBM Enterprise Storage Server, a second-generation Seascape disk storage system, was designed from the ground up to support a new approach to disk storage in networked computing environments — one that is widely accessible, extraordinarily responsive, and centrally managed. With concurrent support for virtually all platforms, advanced performance features, a snap-in design offering easy and quick upgrades, and extremely scalable capacity, the IBM Enterprise Storage Server is one of the most powerful and versatile disk storage solutions available.

While architected for today’s most demanding applications, the IBM Enterprise Storage Server also is designed to easily integrate into your SAN implementations. See product information at:

```
http://www.storage.ibm.com/hardsoft/products/ess/ess.htm
```

3.2.5 LTO

LTO (Linear Tape — Open) is a joint IBM, HP, and Seagate initiative to create new de-facto tape standards in the Open Systems markets. The group developed the formats to serve multiple market areas and to be supported by multiple suppliers. These formats provide customers with a clear and straightforward technology road map, destined to have broad industry support.

Licenses are available to all manufacturers for the two formats based on the technology: Ultrium™, a high-capacity, single-reel implementation that offers up to 200 GB of capacity assuming a 2:1 compression ratio (100 GB native) and Accelis™, a fast-access, dual-reel implementation that offers data retrieval in under 10 seconds.

Magstar technologies are at the foundation of the LTO specifications. This includes the extension of the Magstar method of writing data, the linear serpentine Magstar track recording, an enhanced servo tracking mechanism based on the Magstar MP servo system, a compression scheme derived from the same algorithm as the Magstar LZ1 algorithm, error correction code based on the Magstar architecture, magnetically sensitive highly stable Metal Particle media used in Magstar today, and the implied use of MR heads already in the Magstar products.

The important fact regarding SAN solutions is the native Fibre Channel interface that will be available on products of the LTO format. This enhances configurations and it is not necessary to use a gateway. The device can be directly attached to switches for easier fan-out.

Figure 22 shows the two LTO Tape Formats.

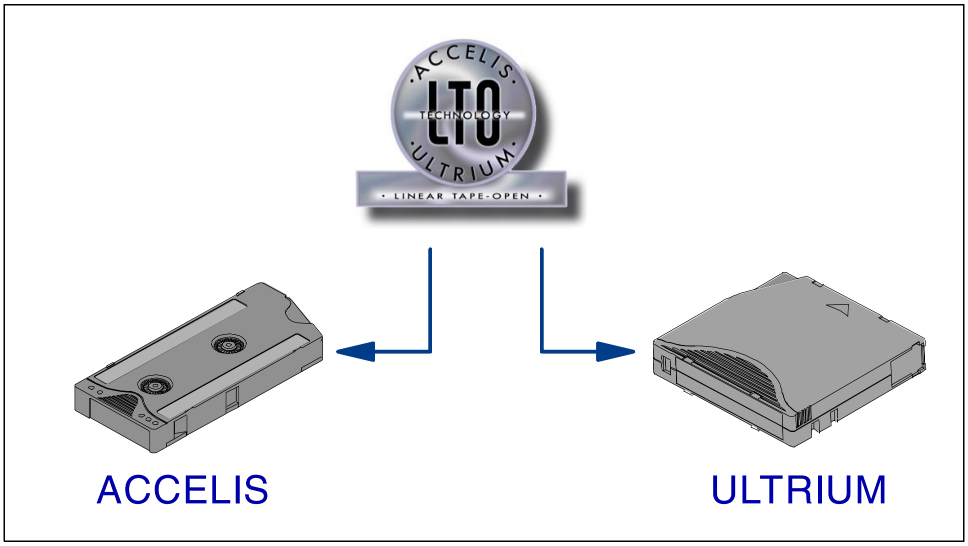

Figure 22. LTO tape formats

Ultrium

The Ultrium tape format is the implementation of LTO technology optimized for high capacity and performance with outstanding reliability, in either a stand-alone or an automated environment. The Ultrium tape format uses a single reel cartridge to maximize capacity. It is ideally suited for backup, restore, and archival applications. They will load in a manner similar to DLT and will fit in DLT automation slots.

Following you will find some specifications of the Ultrium tape format:

- The first generation of Ultrium allows for 100 GB native capacity in a single compact cartridge. The cartridge is smaller than existing single-reel tape cartridges.
- Ultrium provides for data transfer rates of 10-20 MB/sec with the 8-channel head version and 5-10 MB/sec with the 4-channel head version.
- Ultrium provides for 4 different cartridge capacities (10 GB, 30 GB, 50 GB, and 100 GB).
- LTO-CM (Cartridge Memory) enhances functionality by providing a redundant file log as well as user defined information. A nontouching external reader allows immediate access to that information without having to insert the cartridge into a drive.

Accelis

The Accelis tape format is the implementation of LTO technology optimized for fast access to data, with exceptional reliability and performance characteristics. It uses a two-reel cartridge that loads at the middle of the tape to minimize access time. The Accelis tape format is targeted at automated environments and can enable a wide range of “on-line” data inquiry and retrieval applications.

- The first generation of Accelis allows 25 GB native capacity. Self-enclosed tape path in the cartridge eliminates tape threading, which greatly improves time to first data byte.
- Cartridges are loaded in the middle of the tape rather than at the beginning, reducing average search time for random files.
- Accelis provides for data transfer rates of 10-20 MB/sec with the 8-channel head version or 5-10 with MB/sec with the 4-channel head version in the first generation.
- Accelis is ideal for library use, with high-speed access to relatively short files. Accelis is suited for applications like data mining and image retrieval, as well as traditional backup/restore.
- LTO-CM (Cartridge Memory) enhances functionality by providing a redundant file log as well as user-defined information. A nontouching external reader allows immediate access to that information without having to insert the cartridge into a drive.

Refer to these Web sites for technology and product information:

```
http://www.storage.ibm.com/hardsoft/tape/lto/lto.htm
http://www.storage.ibm.com/hardsoft/tape/lto/prod_data/g225-6847.html
```

3.3 IBM SAN Connectivity layer

This section describes IBM SAN connectivity products (presently announced) that belong to the IBM SAN Initiative Connectivity layer (Figure 23) of the IBM SAN Initiative pyramid. These SAN components are not (yet) at a stage where you can safely assume that you can freely use them together and they will work. In fact, the opposite is true: only SAN components that have been tested and verified to work together, from the host to the device, are formally supported.

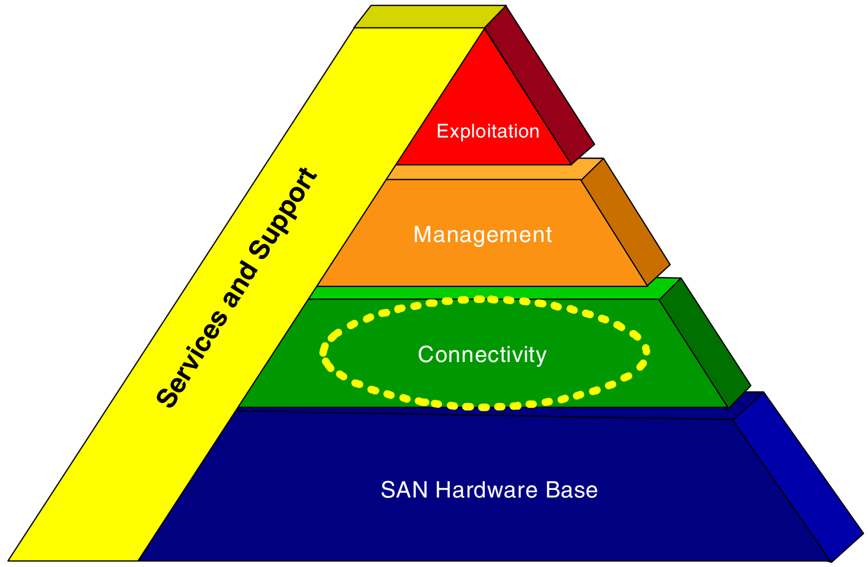

Figure 23. SAN Connectivity layer of IBM SAN Initiative

For this reason, IBM (and all the other companies) will announce and support individual devices only in a precise context, after having validated, in *interoperability labs*, that they work properly as a given set of products. We will indicate, for each product, the contexts in which they are formally supported today.

3.3.1 IBM Fibre Channel Switch

The IBM SAN Fibre Channel Switch offers two models, an 8-port model and a 16-port model (IBM 2108 S08 and S16). It uses a non-blocking switch architecture, thus delivering multiple, concurrent 100 MB/sec connections.

The IBM Fibre Channel Switch offers also advanced management capabilities for:

- Automatic discovery and registration of host and storage devices
- Intelligent rerouting of connection paths, should a port problem occur.

Other features of the IBM Fibre Channel Switch include:

- A Web browser interface that will help you in configuring and managing the switch
- Cascading of switches, for scaling to larger configurations and to provide resiliency for high data availability

- An option for a second power supply
- Configuration with hot pluggable port Gigabit Interface Converters (GBICs) for shortwave or longwave optical connections of up to 10 kilometers

The IBM Fibre Channel Switch supports attachments to multiple host systems:

- IBM Netfinity and Intel-based servers running Microsoft's Windows NT or Novell Netware
- IBM RS/6000 running AIX
- SUN servers running Solaris

The SAN connectivity products and storage systems that can be attached are:

- IBM SAN Data Gateway with IBM Magstar and Magstar MP libraries, and the IBM Versatile Storage Server
- IBM DLT tape libraries
- IBM Fibre Channel Hub and Netfinity Channel Hub
- IBM Fibre Channel RAID Storage Server and the Netfinity Fibre Channel RAID Controller Unit

3.3.2 IBM SAN Data Gateway

The IBM SAN Data Gateway (2108-G07) is one of the first components of the IBM SAN solution that allows an easy migration to the SAN environment using Fibre Channel technology. The SAN Data Gateway connects SCSI and Ultra SCSI storage devices to Fibre Channel environments. It attaches new or existing SCSI storage products to SAN using an industry standard Fibre Channel Arbitrated Loop (FC-AL) interface. The SAN Data Gateway solves three immediate problems:

- The 25m cable length restriction for SCSI — the cable can extend up to 500m
- The increased bandwidth demand the Ultra SCSI storage products can place on the SCSI bus
- The address limitations of SCSI

Hub and switch attachment to the gateway are supported as well. The use of hubs in SAN configurations increases the device connectivity, but hubs have some issues with multiple hosts on the FC-AL loop such as Loop Initialization Process (LIP) and arbitration. If a system is turned OFF/ON or rebooted it might impact the operation of other systems in the FC-AL loop. Many integrators will not support multi-host loop at all.

The use of switches increases the host fan-out which means the number of host connections of SAN configurations.

The SAN Data Gateway is a protocol converter from FC-AL to SCSI that currently supports the attachment of Magstar MP 3570, Magstar MP 3575 and Magstar 3590 to Fibre Channel capable Intel-based servers running Microsoft Windows NT.

IBM supports, or will support in the near future, the following additional product attachment to the SAN Data Gateway:

- Magstar MP 3570 Tape Subsystem and 3575 Tape Library Dataserver attachment to RS/6000 AIX servers, and Sun Solaris servers.
- Magstar 3590 Tape Subsystem in a stand-alone environment; in a Magstar 3494 Tape Library, and Magstar Silo Compatible Tape Subsystem attaching to Intel-based NT servers, RS/6000 AIX servers, and Sun Solaris servers.

It also supports attachment of the IBM Enterprise Storage Server (ESS) to Fibre Channel capable Intel-based servers running Microsoft Windows NT, and Fibre Channel capable UNIX servers from Sun.

You should check the actually supported servers, adapters, disk and tape subsystems on the SAN Data Gateway on:

`http://www.storage.ibm.com/hardsoft/products/sangateway/supserver.htm`

Sharing the gateway between disk and tape products is currently not supported or not practical, because:

- The Enterprise Storage Server needs all the SCSI attachments
- The levels of the HBA (host bus adapter) driver required for disk and for tape are different, which makes it impossible to use gateway-attached disks and tapes on the same host. This will eventually be fixed, but is a nice illustration of an interoperability problem.

The gateway can either be used as a stand-alone table top unit or mounted in a standard 19" rack. The rack can be either the IBM 2101 Seascape Solutions rack or an industry standard rack.

The SAN Data Gateway is equipped with:

- 4 Ultra SCSI Differential ports
- 1 to 3 FC-AL short wave (500 m) ports and Fibre Channel optic cables
- StorWatch SAN Data Gateway Specialist (included on CD)

Features and functions of the SAN Data Gateway

- SAN connectivity:

  Creates reliable SAN solutions without needing hubs, switches and bridges. The SAN Data Gateway provides a distance or connectivity solution for SCSI attached storage devices.

- Heterogeneous systems and storage:

  Provides seamless support for different host platforms and multiple device types.

- SAN resource sharing:

  Zoning or partitioning enables a simple and effective resource sharing solution. Zones are created by controlling the access between different channels or ports and are implemented with the StorWatch SAN Data Gateway Specialist access control function.

- SAN value add functions:
  - Supports up to 256 LUNs across multiple interfaces.
  - Persistent Address Maps are preserved in non volatile memory.
  - Full awareness of SCSI 3 protocol for disk and tape.
  - ’SCSI over TCP’ for remote transfer, management and control. SCSI commands and data are encapsulated in TCP packets.
  - Support for SNIA Extended Copy Command specification; this is the basis for Server-free Backup solutions in the future.

- Transparent SAN performance:

  The total bandwidth of the SAN Data Gateway is 120 MB/sec; the overall performance is driven by the maximum available device performance.

- SAN Management:

  The SAN Data Gateway is remotely managed and controlled by the StorWatch SAN Data Gateway Specialist.

- SAN Scalability:

  The SAN Data Gateway is offered with up to 3 FC ports to provide 3x4 configurations.

Zoning or access control

The SAN Data Gateway has the ability to connect to more than one host. In the default configuration, there is no restriction between the channels for access to the target devices. Without additional controls, host operating systems do not handle multiple systems using the same target devices simultaneously. The result is corrupted file systems when two hosts try to use the same disk drives or LUN. Or, tape backup and restore operations might be interrupted. Use the IBM StorWatch SAN Data Gateway Specialist Channel Access options to disable the desired access between the SAN Connections and individual SCSI channels.

This Access Control function is also called zoning, and it partitions the SAN configuration by either allowing or denying access between the FC and SCSI ports of the gateway.

Figure 24 shows a zoning example for the SAN Data Gateway with two FC ports. In the example we implemented two zones. Zone one with a host on FC adapter 1 sees only the devices attached to SCSI adapter 1 and 2. Zone two contains a different host attached to FC adapter 2 and the devices attached to SCSI adapter 3 and 4.

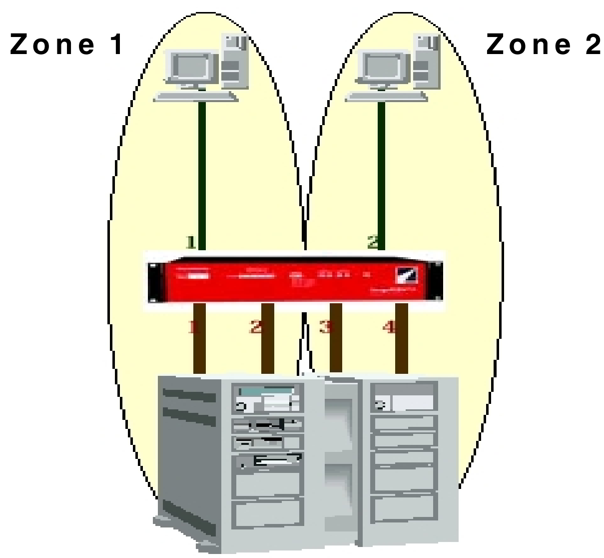

Figure 24. SAN Data Gateway zoning

The SAN Data Gateway supports one server per FC-AL adapter in point-to-point connections. Each host has to have its own drives or tape libraries. This means that there is no drive sharing possible at all. To use a drive on a different host requires that you reconfigure the gateway.

Advantages of SAN Data Gateway versus hub

- Concurrency: aggregate throughput is not limited to one loop.
- Zoning: access control is available based on FC and SCSI ports.
- Hosts are each point-to-point and they can be heterogeneous.
- Smaller configurations with fewer devices offer lower administration cost for customer and lower service/support cost for IBM (as it is easier to isolate problems).
- It avoids the inherent multi-host issues of the FC-AL loop, such as the Loop Initialization Process (LIP) and arbitration. If a system is turned OFF/ON or rebooted, it might impact the operation of other systems in the FC-AL loop. Many integrators will not support multi-host loop at all.

Advantages of SAN Data Gateway versus switch

- It defers or completely avoids the high entry cost of a switch.
- Smaller configurations with fewer devices offer lower administration cost for customer and lower service/support cost for IBM (as it is easier to isolate problems).
- Interoperability issues with switches: fabric support is limited, resource sharing requires middleware.

3.3.3 IBM SAN Data Gateway Router

The SAN Data Gateway Router is a SCSI to Fibre Channel Gateway, with one Fibre Channel adapter and up to two SCSI ports. It is a low-cost solution, compared to the IBM SAN Data Gateway product, which offers up to 3 FC x 4 SCSI ports configurations.

The IBM SAN Data Gateway Router (IBM 2108-R03) can accommodate either UltraSCSI single-ended ports or UltraSCSI differential ports.

The router supports full mapping of SCSI IDs and LUNs between the Fibre Channel attached host and the SCSI tape library. Also, a StorWatch Specialist for centralized configuration and management support is available.

The IBM SAN Data Gateway Router can be attached to an IBM Fibre Channel Switch for extended connectivity.

Host attachment support initially includes:

- IBM Netfinity and Intel-based Windows NT

Support for UNIX and Novell systems is planned for a later date, and specific information on supported configurations will be provided at a later date.

3.3.4 IBM Fibre Channel Storage Hub

The IBM Fibre Channel Storage Hub (2103-H07) was announced on February 16, 1999, as part of the IBM Fibre Channel RAID Storage Server announcement. See the specification sheet at:

```
http://www.storage.ibm.com/hardsoft/products/fchub/fchub.htm
```

It is also part of the IBM Netfinity Fibre Channel Storage Solutions.

The IBM Fibre Channel Storage Hub is a 7-port central interconnection for Fibre channel Arbitrated Loop (FC-AL) that follows the ANSI FC-AL Standard. Each Fibre Channel Storage Hub port receives serial data from an attached node and retransmits the data out of the next hub port to the next node attached in the loop. Each reception includes data regeneration (both signal timing and amplitude), supporting full-distance optical links.

The Fibre Channel Storage Hub detects any loop node that is missing or is inoperative and automatically routes the data to the next operation port and attached node in the loop.

Each port requires a Gigabit Interface Converter (GBIC) to connect to each attached node. The Fibre Channel Storage Hub supports any combination of short-wave or long-wave optical GBIC’s. The GBICs are hot-pluggable into the Fibre Channel Storage Hub, which means that you can add servers and storage devices to the arbitrated loop dynamically, without powering off the Fibre Channel Storage Hub or any connected devices. If you remove a GBIC from a Fibre Channel Storage Hub, that port is automatically bypassed. Conversely, if you plug a GBIC into the Fibre Channel Storage Hub, it will automatically be inserted and become a node on the loop.

Hubs can also be used as distance extenders, in connection with the SAN Data Gateway, as shown in Figure 25 on page 58. Note that this is an example of using hubs as distance extenders, but this is not an IBM configuration at this date.

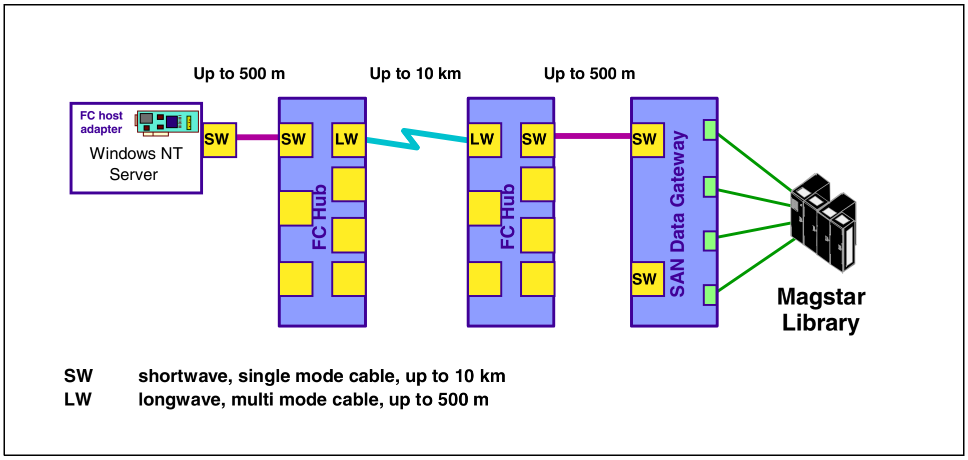

Figure 25. Using hubs for greater distances

Today, the IBM Fibre Channel Hub is only supported in an NT environment.

3.3.5 Fibre Channel Adapters

Fibre Channel adapters, also called Host Bus Adapters (HBAs), are used to attach host systems to a Fibre Channel network. Several variations are possible on:

- Type of cable supported: long-wave fiber, short-wave fiber, copper
- Fibre Topology support: point-to-point, loop, switch attachment capabilities
- Protocol support: FCP-SCSI, IP, and SCSI concurrently; other FC-4 protocols: IPI, SB-2

IBM has announced two FC adapters:

- The Netfinity Fibre Channel PCI Adapter, which is in fact the Qlogic 2100F adapter.
- The RS/6000 (TM) Giga Fibre Channel Adapter for PCI Bus (FC #6227), which is the Emulex LP7000 adapter.
- IBM also has announced, as part of the IBM Fibre Channel RAID Storage Server announcement, support for FC Adapters on SUN and HP UNIX servers. See this announcement for detailed technical information.
- The Qlogic 2100F adapter is also available for non-Netfinity Windows NT hosts.

You can reach the Qlogic Web page at:

`http://www.qlc.com/`

You can reach the Emulex Web page at:

`http://www.emulex.com/`

These two FC adapters have the characteristics shown in Table 1.

Table 1. FC adapter characteristics

Adapter	Topologies	Cabling	Protocols
RS/6000 Giga Fibre Adapter for PCI	Point-to-point Arbitrated loop Switch fabric	Fiber Copper	SCSI FTP
Netfinity Fibre Channel PCI Adapter	Arbitrated loop	Fiber	SCSI

3.4 IBM SAN Management layer

The SAN Management layer (Figure 26) of the IBM SAN Initiative pyramid provides comprehensive control of all storage resources through a centralized focal point. This leads to reduced costs due to an effective and flexible management of the storage resources.

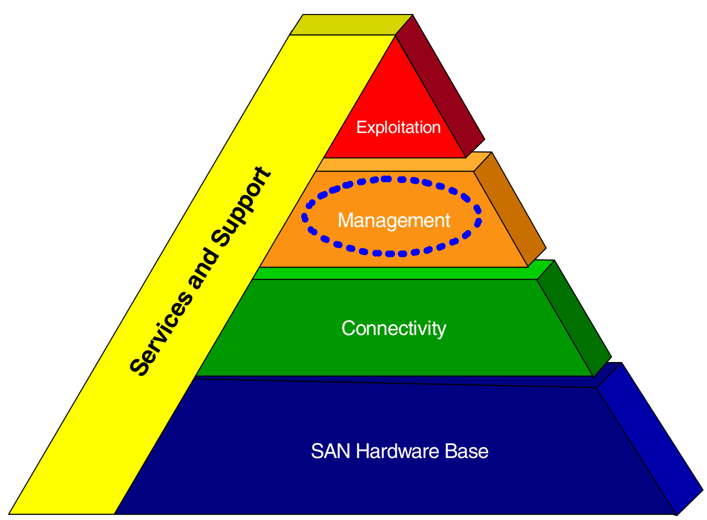

Figure 26. SAN Management layer of IBM SAN Initiative

3.4.1 IBM SAN Management Initiative software

Software for the SAN Management Initiative is needed to configure, control and monitor the SAN and all of its components in a consistent manner. This software may include management interfaces to individual devices, but ultimately will present a single interface to manage all SAN resources.

Currently there is no single generalized tool available to manage all the SAN components, but there are already pieces available. Those approaches can be categorized as follows:

- **Device level management**

  Device level management includes the configuration, monitoring and simple management of individual Fibre Channel components such as switches, hubs, and gateways but also individual devices like the IBM Enterprise Storage Server (ESS).

  Examples for device level management products are the IBM products StorWatch SAN Data Gateway Specialist, the StorWatch Enterprise Storage Specialist and the StorWatch Fibre Channel RAID Specialist.

  Examples for switch management tools are the following Brocade software tools: Brocade SES (SCSI-3 Enclosure Services) and Brocade Web Tools.

- **Total SAN management or SAN level management tools**

  No total SAN management tools are available so far.

Most of these tools work with GUI interfaces and gather the required information through either:

- Inband interface:

  This is the management of SAN components over the data path through SCSI protocol.

- Outband interface:

  Required data is gathered, for example, through an Ethernet or RS-232 serial interface from the SAN component.

IBM approaches a total SAN management solution in the form of a new StorWatch expert component which will be available in 2000. It has a collection of tightly coupled features commonly associated with storage management. The features all rely on a core set of functions provided by the program infrastructure and a central database of information collected and analyzed by each feature.

The configuration and management done for the SAN is supplemental to, or an extension of, the configuration and management capabilities of the participating operating systems. In other words, IBM's SAN management solution utilizes operating system (and FC component) functions as appropriate to implement its functions. The host view, switch view and storage device views thus become part of a new, higher, SAN view.

The database is a repository for information about the SAN and its components. The information includes but is not limited to device Vital Product Data, physical configuration, logical configuration, interconnects, state, historical performance, and service data.

SAN level management involves management of the entire switched fabric (and its components and storage devices) as a whole. There are many aspects of total SAN management. Some of the simpler, lower value functions include:

- Device installation, configuration and monitoring
- Keeping an inventory of resources on the SAN
- Reporting utilities
- Assistance with SAN planning

Some of the more advanced, higher value functions include:

- Automated component and fabric topology discovery
- Managing the fabric configuration including zoning configurations
- Name services
- Managing security
- Health and performance monitoring
- Workload management/balancing of SAN activities based on application priority (for example, backups are lower priority than business transactions)

As you can see, SAN management functions differ little from traditional systems and storage management functions. As such, IBM's SAN management solutions are integrated well with IBM's systems management solutions (Tivoli) and storage management solutions (DFSMS, Tivoli Storage Manager, StorWatch).

You can imagine a total SAN management solution as a SAN Manager available as a further component in the Tivoli software family. This Tivoli SAN Manager would work together with StorWatch specialist agents on the managed resources. Figure 27 shows a possible total SAN management approach.

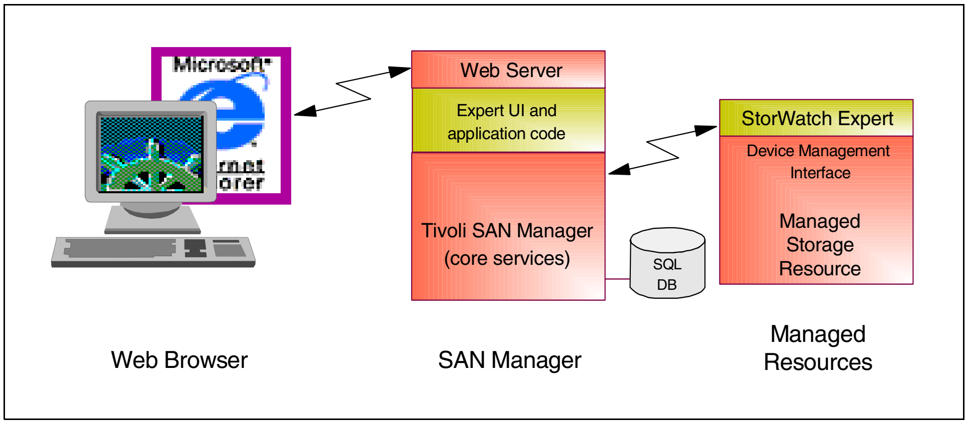

Figure 27. Total SAN management approach

3.4.2 IBM StorWatch SAN Data Gateway Specialist

This section is an introduction to the StorWatch SAN Data Gateway Specialist software, which is working on the device level management in the SAN management initiative. See 3.4.1, “IBM SAN Management Initiative software” on page 60 for an explanation of the device level management.

The StorWatch SAN Data Gateway Specialist is a client/server application that helps to manage the SAN Data Gateway with different functions. The SAN Data Gateway can be implemented and managed also without this software by using the SAN Data Gateway Service Port, but the Specialist allows a central management of all gateways in an installation through a GUI interfaces and is more user-friendly. The Specialist implements the following functions:

- Asset management:
  - Descriptions and product data for gateways, channels, devices
  - View of current SAN Data Gateway front panel status
- Configuration management:
  - Determine current topology; set gateway, channel, and device parameters
  - Upgrade microcode of gateways from central point

- Device and event monitoring:
  - Health checks, set SNMP alert thresholds, eventlog
  - Error analysis to isolate Fibre Channel and SCSI problems

The IBM StorWatch SAN Data Gateway Specialist can be used to remotely configure and monitor multiple SAN Data Gateways. The Specialist uses a combination of industry-standard SNMP (Simple Network Management Protocol) requests for event monitoring, and a method or technology, known as SCSI, over TCP, which encapsulates SCSI commands and/or data in TCP packets for gateway management.

Client/Server model

Figure 28 shows how you can implement the IBM StorWatch SAN Data Gateway Specialist as a three-part Client/Server model.

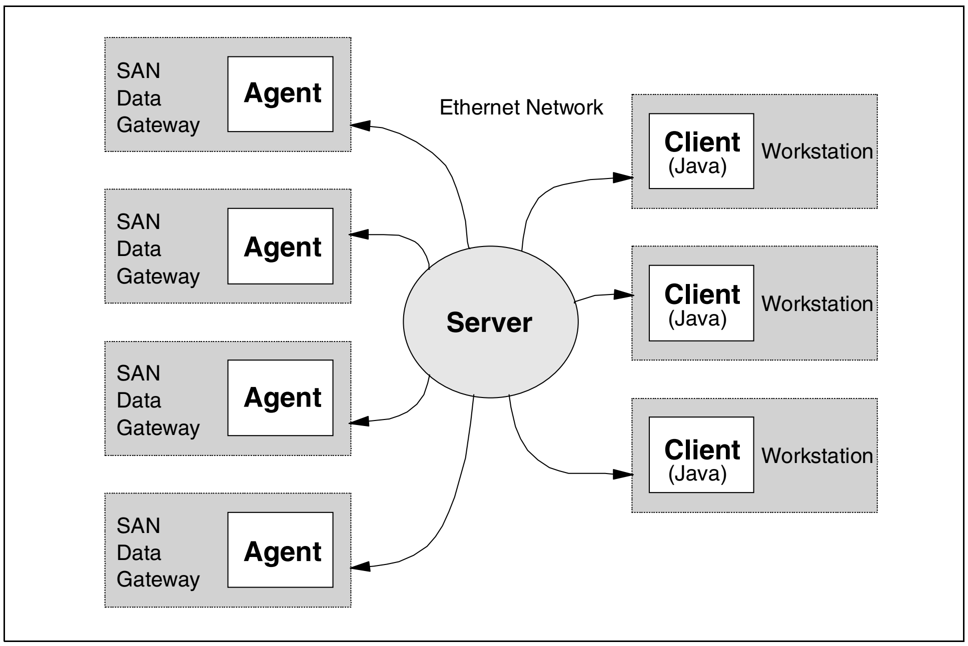

Figure 28. SAN Data Gateway Specialist application model

Agent

Each gateway is a stand-alone, SNMP-manageable host. An agent component is embedded in the operating software of the gateway. The IBM StorWatch SAN Data Gateway Specialist uses SNMP as the primary method of communication with the agent. It allows you to set and retrieve information that control the operation of the agent, to update the gateways and target devices and manipulate device operating parameters. It also provides alerts (traps) when an event has occurred that requires intervention.

Server

The Server component is a Java application that runs on a Windows NT system. In a simple configuration; the Server and Client can be on the same host. It is responsible for maintaining communication with the managed agents, and acts as an intermediary between the Agent running on the gateway and multiple Clients. It provides security features by maintaining account names and passwords on behalf of the Client application. By keeping track of different Client configurations, a user can recall a saved configuration from any Client.

The server coordinates the request from multiple clients to manage multiple gateways. All communications between the clients and the server is implemented using the Java Remote Method Invocation (RMI), a form of Remote Procedure Call (RPC). The server is written to be both transportable and efficient. Multiple clients can share data the server already knows about rather than having to request the data again. In addition, the server receives all traps from the Agents and forwards them to the clients that have been registered to receive them.

Client

The Client is also a Java application. One or more Clients connect to a Server in order to manage the gateways. The Client operates from any compatible computer, as long as a Transmission Control Protocol/Internet Protocol (TCP/IP) connection is established to the Server component. This allows for dial-in configurations using Point-to-Point Protocol (PPP), Intranet and Internet access (where allowed by local network policy and firewall configurations).

The Client is the application that provides the user interface and allows viewing and manipulating all gateway and device parameters. Each Client can be configured by the individual user to show only the gateways of interest. This means that one Client can monitor one set of gateways and other managers can be responsible for totally different gateways, without interfering with each other.

The Client uses Java RMI calls to communicate with the server and SCSI over TCP to communicate to the gateways for some operations.

General implementation considerations

The StorWatch SAN Data Gateway Specialist is a Windows NT based Client/Server application, running on Windows NT 4.0 Server or Workstation, Service Pack 4 or later. The Windows NT systems for both the Server and the Client component require an Ethernet connection with TCP/IP protocol installed. This is required to gather the required data from the gateway. In addition to the application components, the Java Runtime Environment (JRE) is installed.

The StorWatch SAN Data Gateway Specialist can be either installed from the CD delivered with the SAN Data Gateway or by downloading the software through the SAN Data Gateway Internet page:

`http://www.storage.ibm.com/SANgateway/`

In any case, you should check the Web site for the latest recommendations for Operating System, Firmware and the StorWatch SAN Data Gateway Specialist software.

3.4.3 Zoning or access control

Zoning enables resource sharing SAN solutions, but sharing in this case does not mean having shared access to the same drives. It is related to the partitioning of resources. An example is the IBM 3575 Magstar MP Tape Library. One logical library in the IBM 3575 could be attached through the SAN Data Gateway to one server, and the second logical library could be attached through the SAN Data Gateway to a different server. The status can be changed by reconfiguring the SAN Data Gateway, but this is not really dynamic. The gateway provides additional flexibility in library partitioning, since there is no requirement for recabling, but only for reconfiguration.

Table 2 helps with the understanding of tape library or device sharing / partitioning requirements. Tape Library sharing we consider as partitioning, where as the tape drive sharing is real sharing, having dynamically access to a set of drives from multiple hosts.

Table 2. Library /drive partitioning versus sharing

	Tape library sharing	Tape drive sharing
Shared by homogeneous applications	Solved by backup application (ARCserve, Tivoli Storage Manager*) OR by tape library hardware partitioning (IBM 3570, 3575, 3494)	Solved by backup application (ARCserve, Tivoli Storage Manager)
Shared by heterogeneous applications	Solved by Library Management Middleware OR by tape library hardware partitioning (IBM 3570, 3575, 3494)	Solved by Library Management Middleware (Veritas Media Librarian, Legato SmartMedia, Tivoli Removable Media Manager)

Zoning is a crucial function for SANs to support. Zoning is required for a number of reasons. Some operating systems like Windows NT assume that they own all devices that they can access; therefore, they must be restricted from accessing devices that are not intended to belong to them. Prevention of access to selected storage nodes from selected hosts may be necessary for reasons such as security, separation of test from production, and contractual obligations for isolation.

3.5 IBM SAN Exploitation layer

The SAN Exploitation layer (Figure 29) at the top of the IBM SAN Initiative pyramid includes the products and functions that will allow you to make good use of a SAN. In fact, the availability of these functions, and the value these can bring to your enterprise, is certainly the reason for your decision to install a SAN at all.

We start with a section on Tivoli Storage Manager, in which we explain to you how a backup/restore product can make use of a SAN. Some of these considerations apply to any tape application, including other backup products.

The other theme we review is a generalized tape sharing solution, its goal being to allow any application in any environment to access your tape devices and libraries on the SAN.

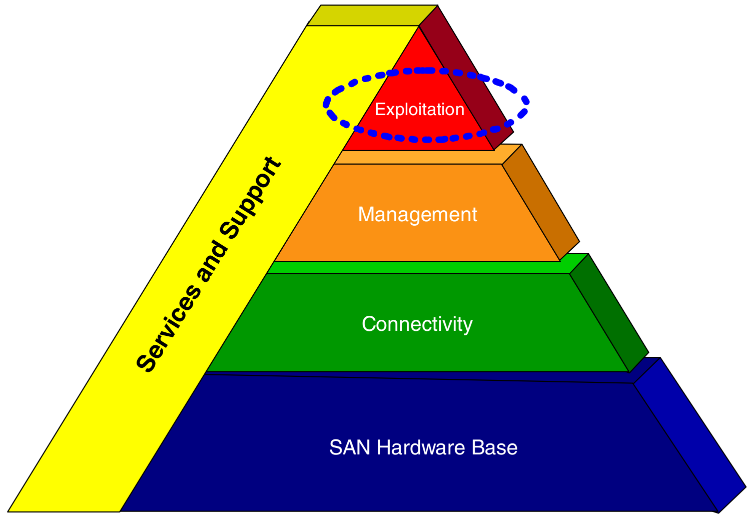

Figure 29. SAN Exploitation layer of IBM SAN Initiative

3.5.1 Tivoli Storage Manager

Tivoli Storage Manager is the follow-on product to ADSM.

Tivoli Storage Manager can already take advantage of SAN attached disk and tape devices, as the SAN is transparent to the applications (storage devices being viewed as just ordinary SCSI-attached devices), even if some additional installation considerations apply; for example, to define tape drives attached across a SAN Data Gateway. These are documented in the Device Readme file for the ADSM NT server at level 3.1.2.20.

But to really exploit the new opportunities offered by SAN technologies, there are some changes, occasional redesigns, and additional functions that will have to be made to Tivoli Storage Manager. We review here some of the changes that we should see in the near to midterm future. Some of these functions have already been mentioned, while some of them are more speculative and merely hint at what it would be possible to do, without any implication that these will actually appear in the product in that form.

Please refer to the Tivoli Web site for current product specifications:

```
http://www.tivoli.com/products/index/storage_mgr/
```

Having Tivoli Storage Manager in a SAN environment provides some interesting new possibilities (refer to 2.2, “New opportunities” on page 26 for an overview of these functions):

- One special form of “LAN-free” client consists in using the SAN for the TCP/IP traffic between the Client and the Server, as shown in Figure 30. This allows you to take advantage of increased throughput (100 MB/sec) compared to classical LANs, and separates the backup traffic from your production applications using the LAN. In a switched SAN, each session will have the full Fibre Channel bandwidth available. Note that the load on the client and on the server nodes does not decrease in this scenario, it even goes up with the increased throughput, as you still need to go through the TCP/IP protocol stack on each host.

  This function requires support for TCP/IP host-to-host communications over Fibre Channel and supported adapter cards.

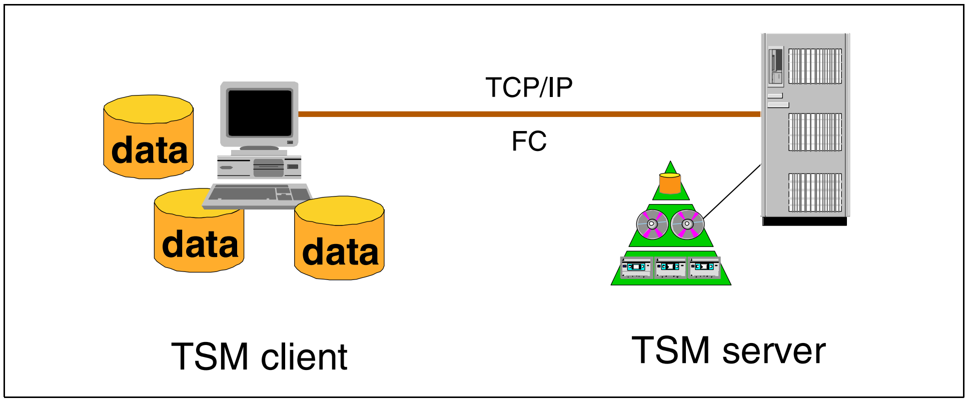

Figure 30. Client-server transfer over Fibre Channel

- Extended distance for tape devices: this allows you to place your tape drives outside of the machine room, away from the server. You can make use of this for your Primary tape volume and/or for your disaster recovery volumes (electronic vaulting). This does not require any additional function in Tivoli Storage Manager.
- Extended distance for disk units: by having disk units in a remote location, you can place a Tivoli Storage Manager mirror copy of your database and recovery log on a disk in a remote location, thus allowing you to restart Tivoli Storage Manager services very quickly after a disaster at the primary site. This does not require changes to Tivoli Storage Manager.

- Direct client-to-Tivoli Storage Manager storage pool volume backups and restores: a client could send its data directly to a Tivoli Storage Manager storage pool volume, either disk or tape, under control of the Tivoli Storage Manager server who would tell the client where (to which disk sectors, for example) to write or read that data on the storage pool volume. This is often referred to as LAN-free backup, because the data does not flow through the customer LAN.

  The client session initialization and control information flows over LAN, using TCP/IP protocols, but file data flows directly over the SAN, using SCSI protocols. The big gain here is a decrease of the load on the Tivoli Storage Manager server. This requires new functions in both the client and the server.

- Third-party copy services: SANs allow for third-party copy functions, where data goes directly from one storage device (for instance, a disk) to another one (disk or tape) without involving any host for the data transfer. The advantages of using these copy services are a reduction in host overhead and, possibly, improved performance.

  For Tivoli Storage Manager functions, this could be used for:

  - Server processes, namely storagepool migration, storagepool backup, and tape reclamation, which would bring a significant reduction of the load on the Tivoli Storage Manager server host.
  - Client Sessions: backup, archive, restore, and retrieve traffic would go directly from the client disk (on the SAN) to server disks or tapes (also on the SAN). This would reduce the load on both the client and the server hosts. This will be of interest only for client data that reside on the SAN attached disk.

  Direct client-to-server transfers open also new opportunities for client functions, like taking a full volume backup of a disk drive without impacting the client node at all.

  Use of third-party copy services will clearly require changes in client and server code.

- Tape connectivity: being able to connect tape drives to multiple hosts is a big plus if you have more than one Tivoli Storage Manager server host. It allows you to have essentially as many Tivoli Storage Manager server as you need, without needing to install drives on each server machine. Reasons for using multiple additional Tivoli Storage Manager server hosts are:

- Load balancing: having two smaller server instances instead of one large server allows you to keep the load of the server machine at a comfortable level and/or to keep the Tivoli Storage Manager database size to a reasonable amount, all of which helps performance.
- Separate backup and archive servers: separating backup activities, where you typically keep backup versions for a relatively short period of time (a few weeks) and archive activities, where the retention periods are normally years, can be very effective in optimizing the Tivoli Storage Manager database activities, and in making it easier to treat these two types of data differently, for example, regarding disaster recovery.
- Dedicated Tivoli Storage Manager servers: you may need, for instance for security reasons, a separate Tivoli Storage Manager server for a set of clients or for a specific application. You can now do this, even multiple times, without necessarily purchasing additional tape drives.
- Validation servers: many installations have implemented change management policies that require any change to be tested on a validation system before being installed in the production environment. For validating your Tivoli Storage Manager server, you will need access to tape drives and libraries, and, thanks to SAN extended connectivity, you will be able to do this without requiring additional equipment, given that the tape drive requirements of a validation server are limited in time.
- Server take-over: if you need a stand-by machine for your Tivoli Storage Manager server(s), you need also connectivity to the tape library. The extended connectivity you get with SAN allows you to do this without installing additional drives
- One server per (large) client: for clients having very large databases, SAN tape connectivity and pooling make it possible to have a Tivoli Storage Manager server on the same machine as the database, using any tape drive in the SAN. This is a special case of the dedicated server case, mentioned here for performance reasons.

If you already have several Tivoli Storage Manager servers, they probably have each a set of tape drives. In this situation, it can happen that a server needs, at a given time, more drives than it has, while another server's drives are idle. Enhanced connectivity allows to use all your tape drives as one pool, thus obtaining a better overall utilization, and/or requiring a smaller number of tape drives. And if you need fewer tape drives, you may then choose to take the best technology, rather than having to settle for (relatively) cheap units just because you need many of them.

Connectivity of multiple Tivoli Storage Manager servers to tape drives and libraries implies that something has to manage tape drive assignments to make sure that no two servers try to use a given tape drive at the same time. This function can be provided either by Tivoli Storage Manager, to allow drive and library sharing between Tivoli Storage Manager servers, or by another, more generalized product, that would provide tape sharing to any application, thus allowing to use the same tape pool for Tivoli Storage Manager and other tape applications.

Unless that tape sharing function is available, you will need to manage the assignments of drives to different servers yourself. For the cases where changes of configurations are not very frequent, this can be done by simply putting the given drive off-line to all the Tivoli Storage Manager servers except for the one needing the drive. This would be the case, for instance, with a validation server.

Using Tivoli Storage Manager with the SAN Data Gateway

We now consider a sample configuration of Tivoli Storage Manager servers and Magstar 3590 drives in an IBM 3494 library, and look at how connecting the tape devices through SAN and the IBM SAN Data Gateway gives you additional flexibility and better resource utilization. These improvements would be even greater if considering tape devices other than Magstar 3590, with only one SCSI port per drive. And of course, these same advantages will remain if the tape drives have native Fibre Channel attachments.

The SAN Data Gateway connectivity capabilities allow you to access up to three hosts from each drive port, or up to six hosts per drive if you use two SAN Data Gateways, instead of only two host connections per drive.

It also gives you better flexibility for the physical installation of your drives, as you now have:

- Extended distance: 500 meters versus 25 meters (without extenders)
- Attachment of drives residing on several SCSI busses through one host FC adapter card, instead of requiring that many SCSI adapters on the host

You can take advantage of these new opportunities by connecting all your Tivoli Storage Manager servers to all the drives installed in your IBM 3494. This allows you to exploit the library drives as one large device pool, instead of partitioning the drives over the Tivoli Storage Manager servers, as you need to do in a classical SCSI set-up, where you cannot connect a 3590 drive to more than two systems. (See Figure 31.)

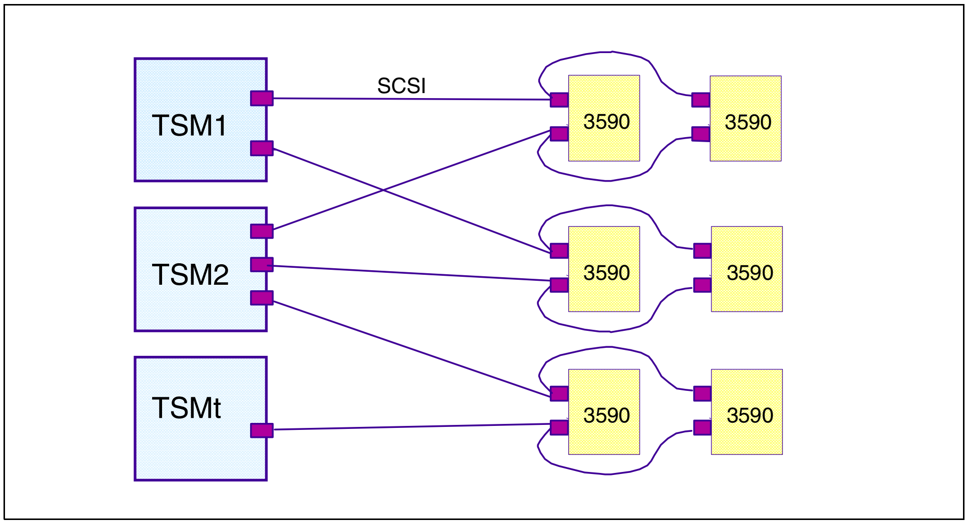

Figure 31. 3590 SCSI connectivity

Compare these two configurations: Figure 31 is typical of an installation that has two Tivoli Storage Manager server production systems, Tivoli Storage Manager1 and Tivoli Storage Manager2, and one test Tivoli Storage Manager server, Tivoli Storage ManagerT, running on another system, and used as a validation and test machine for new software levels or for testing procedures before they go on the production servers. In addition, Tivoli Storage Manager1 and Tivoli Storage Manager2 are fail-over systems for each other.

In this situation, the connectivity restrictions mean that:

- You cannot connect more than two systems to one given 3590 drive.
- The maximum distance to the last drive on the SCSI bus is 25 meters. Given that you need a cable of 3.4 meters inside the 3494 library to daisy-chain the two drives, this means in practice that you need one SCSI adapter in the host per pair of 3590's.
- Your IBM 3494 library can be at most 18 meters away from either of your servers.

The consequences of these restrictions are:

1. Even though your tape load could be satisfied with 4 drives (for example), you need two additional drives to have drives reachable from Tivoli Storage ManagerT.

2. You need two SCSI adapters on Tivoli Storage Manager1, and three on Tivoli Storage Manager2, not because of performance constraints, but only to get the required connectivity on all your servers.
3. You don’t have the liberty of placing your library outside of your machine room because of the distance limitations.

Now compare the same situation with a solution using SAN elements, namely host FC adapters and the SAN Data Gateway, as shown in Figure 32.

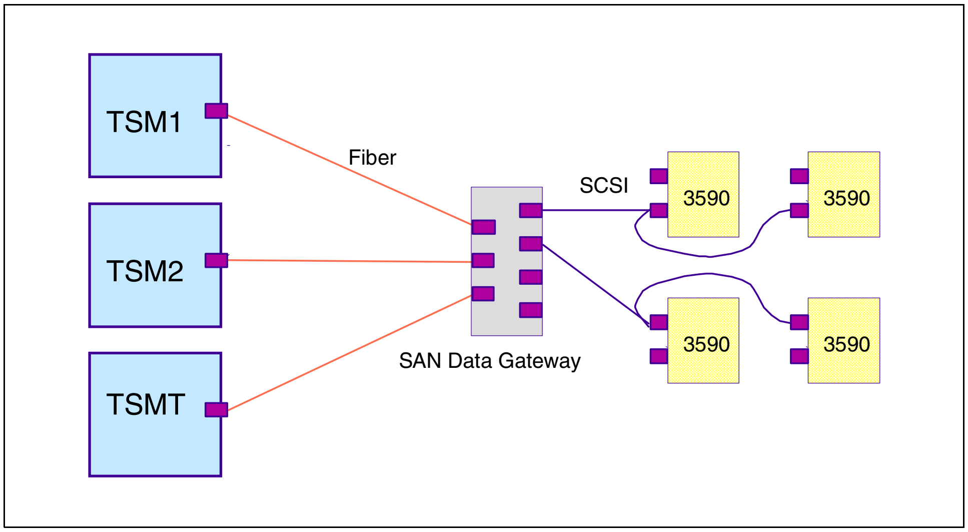

Figure 32. SAN 3590 connectivity

One Fibre Channel adapter in each host gives you connection to all the 3590 drives in the 3494 library. The diagram is symmetrical, all the hosts can see all the drives. The main advantages of this configuration are:

1. The maximum distance between the hosts and the library is now 500 meters, which allows you to place the 3494 library in another building, if you wish to do so for disaster recovery reasons. (Note that having your tapes away from your machine room does not protect you from a disaster in your tape vault: you still need to store the necessary disaster copies of your tapes on another site. Thanks to SAN, these COPY volumes can also be copied across the network, thus avoiding physical handling of the tape volumes).
2. Assuming that your workload does not require more than four 3590 drives, you don’t need to purchase additional drives just for connectivity reasons.

Note that these connectivity possibilities of SAN are necessary, but not sufficient to utilize your tape drives as one resource pool. What you still require is a way to *manage the drive allocation*, in order to avoid the situation of two hosts trying to access the same Magstar 3590 drive at the same time. You also need a Tape Management System (TMS) function to control the allocation of tape volumes (media!) between the different tape servers.

This any-to-any server-to-tape drive connectivity will be of interest in these typical scenarios:

1. Tape server take-over
2. Tape drive sharing
3. Electronic vaulting

Tivoli Storage Manager tape sharing situation

Tape sharing between two hosts, in the sense of having two Tivoli Storage Manager servers on two hosts accessing the same tape drive, has been possible since the availability of the IBM 3590 drive, with its two SCSI ports. However, the only tape sharing configuration supported by Tivoli Storage Manager presently is for two Tivoli Storage Manager servers running on the same machine, as shown in Figure 33 on page 76. The reason tape sharing between two Tivoli Storage Manager servers on different machines is not supported by Tivoli Storage Manager is that concurrent accesses from both hosts are not guaranteed to be foolproof in all cases. The problem is that a RESERVE command (used by Tivoli Storage Manager) does not currently operate across systems. Thus, a Tivoli Storage Manager server cannot be sure that it has exclusive rights to a drive. For example, one Tivoli Storage Manager server could be using one of its tape drives, while the other Tivoli Storage Manager server requests a scratch mount. The tape in the drive will look like a valid scratch tape to the second server.

A new function in the IBM 3590 microcode called auto-share feature is solving that situation, by assuring that RESERVEs coming from different systems are observed. This means that tape sharing by multiple Tivoli Storage Manager hosts is now possible.

Another possibility to have tape sharing between Tivoli Storage Manager server machines is to use the External Library Function of Tivoli Storage Manager, where the actual library control functions (mount, dismount, eject, insert) are managed outside of Tivoli Storage Manager. This Library management function has to be realized by user-written code.

Tivoli Storage Manager tape pooling solution

Tivoli Storage Manager provides a library, drive, and media sharing solution for Tivoli Storage Manager servers. While this does not provide resource sharing between heterogeneous applications, because only Tivoli Storage Manager servers can participate, it is a very common need, and so is very useful in many Tivoli Storage Manager installations.

The Tivoli Storage Manager solution consists in having one Tivoli Storage Manager server acting as the Tape Library server, controlling the tape libraries and drives; the other Tivoli Storage Manager servers are library clients; they will simply send their tape mount requests to that library server in a way similar to what is done today when using the External Library Manager interface.

Tivoli Storage Manager Version 3.7 exploits the Storage Area Network (SAN) technology for dynamic tape library sharing. With this function, multiple Tivoli Storage Manager servers can dynamically share library volume and tape drive resources of one physical tape library. Using the Fibre Channel technology of SANs, distances between the tape library and the Tivoli Storage Manager server can be extended.

One Tivoli Storage Manager server, known as the library server, controls the tape library, which means it is responsible for the mounts and dismounts of tapes into the drives, for the serialization of the drive accesses, and for the library volume inventory. The server can also use all of the tape drives for its own operation.

All other Tivoli Storage Manager servers using this tape library are known as library clients. They have physical access to all the drives throughout the SAN. They can directly read or write data to or from tapes, but they need the library manager server to get tapes mounted or unmounted. At the present time, SAN tape library sharing is only available for NT and Sun servers and for selected types of tape libraries.

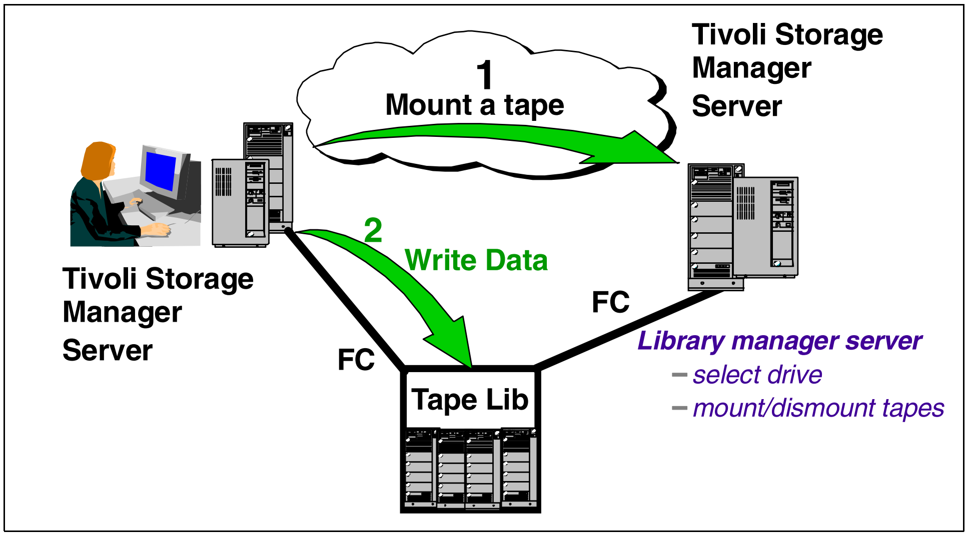

Figure 33. Tivoli storage manager tape sharing

The Library Server is also able to manage one common scratch volume pool for all the connected servers.

This enables the utilization of all your existing libraries and drives as one resource pool, thus bringing improved resource utilization, and it simplifies the management of drives and tape volumes.

Often, these new facilities also bring with them new problems that you will need to address, namely:

- Drive preemption: this is a function in a Tivoli Storage Manager server whereby tape operations have different priorities, and a process of low priority will be cancelled by the server for a process or session of higher priority. For example, a restore session will take precedence over a tape reclaim process. In a tape pooling situation, there is no mechanism for one server to preempt drives used by another server. The only solution will then be to implement that function by external intervention, either human or programmed.
- Two-drive mount requests: one individual Tivoli Storage Manager server will not attempt to mount a volume for a two-drive operation (tape reclamation, tape storagepool backup, moving data from tape to tape) if it does not have two tape mount points available. In a tape pooling situation, the mount requests will be sent one after the other to the Library Manager, without indication that these two requests are related and that mounting just one of the two volumes is not enough. This means that you could get

into deadlock situations, where two servers each hold one drive (in a simple case where you have two drives in total), and each of these servers would wait for a second drive. Again, the only remedy in such a case will be external intervention.

In summary, even though you will get better resource utilization with Tivoli Storage Manager tape pooling, you will still need to ensure that you have enough tape drives to absorb your workload, and be prepared to take corrective actions for the cases where demand exceeds supply.

3.5.2 Tivoli Removable Media Manager

In an Open Systems environment today, you cannot easily share tape devices across multiple hosts, essentially because you don’t have the necessary connectivity between the host systems and the tape drives. SAN provides this connectivity, in principle allowing any host to access any tape drive. This will allow you to treat all you tape devices and tape libraries as one common resource pool, provided you have a management tool to manage these resources.

Tivoli Removable Media Manager belongs to the IBM SAN Exploitation initiative, and is intended to provide a generalized facility for managing tape libraries, tape drives, and tape media from multiple, heterogeneous hosts and applications in a SAN environment.

Tivoli Removable Media Manager will be a client-server application. Typically, Tivoli Removable Media Manager clients (host systems) will ask the Tivoli Removable Media Manager server to mount a tape volume. The Tivoli Removable Media Manager server will be responsible for checking that this client is authorized to access that tape volume, mount it in a suitable drive, and tell the client where it can access the volume.

The Tivoli Removable Media Manager clients will use the Microsoft RSM (Removable Storage Manager) API, widely accepted in the industry, which will facilitate the integration of non-IBM applications.

The server will be responsible for:

- Library management and sharing
- Drive management and sharing
- Media management, including media pools management.

Tivoli Removable Media Manager will ultimately provide these functions at the Enterprise level, allowing you to manage all the libraries, drives and media as one pool of resources, accessible from NT, UNIX, and mainframe applications.

3.6 IBM SAN Services and Support layer

Perhaps the most important part of the IBM SAN Initiative is the promise of providing SAN Services and Support (Figure 34). This layer of the pyramid addresses two main aspects: interoperability, which IBM is conducting to ensure that the many elements found in a SAN will provide you with a guaranteed working solution, and the services to help you design, plan, implement, and operate a SAN in an effective manner.

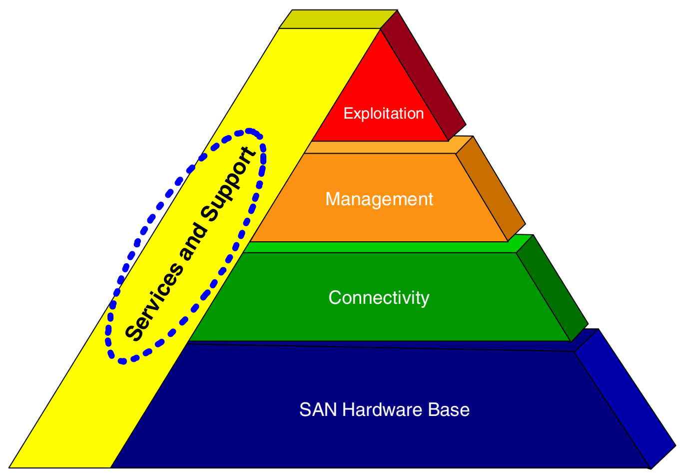

Figure 34. SAN Services and Support layer of IBM SAN Initiative

3.6.1 Interoperability

In this section, we will explain why interoperability issues are important in SAN environments, their impact on product availability, and how IBM is dealing with these issues.

Why any-to-any is not "anything goes"

The Fibre Channel standards specifications, like those of most standards, leave room for interpretation and implementation options. This has the unfortunate consequence that not just any SAN device will work with any other SAN device. All the elements involved in a SAN (FC adapters hardware, microcode, driver, host operating system, cabling, SAN topologies and services, gateways, bridges, hubs, device adapters and their microcode and drivers) must be compatible with each other in order to have working end-to-end communications. So far, given the options left open by the standards, the only way to ensure that a given set of SAN products work together is to try to test them.

You, as a customer, will not expect to buy equipment and perform such tests, only to find out that you must eventually change a number of these elements yourself. Rather, you will expect that your supplier(s) will provide you with a working solution, using the products that they propose to sell you.

This is the reason why IBM will only support sets of SAN products which have been verified to be working properly together. This is established after a series of comprehensive tests conducted in IBM interoperability laboratories.

Given that the number of possible permutations of all the elements involved in a SAN quickly goes to astronomical numbers, IBM will select relevant elements to test, and will announce them by choosing elements for SAN solutions, for example, tape pooling with NT servers.

As testing activities progress, you can expect to see more tested solution scenarios, involving more diverse types of elements.

In addition to these pre-determined solutions, customers may want to test, before purchase, their own sets of SAN elements. For this purpose, IBM will provide Certification Labs, one in Gaithersburg, Maryland, and one in Mainz, Germany, where these tests will be conducted for specific customer requirements.

Interoperability Matrix

You can find a list of hosts, adapters, and tape devices supported with the IBM SAN Data Gateway and the IBM SAN Fibre Channel Switch at:

```
http://www.storage.ibm.com/hardsoft/products/tape/supserver.htm
```

Applications supported in the context of tape devices connected to Windows NT systems using the SAN Data Gateway are listed in the Independent Software Vendor (ISV) matrix available at:

`http://www.storage.ibm.com/hardsoft/tape/`

The ISV matrix currently contains, for the Windows NT platform:

- IBM Magstar MP 3570
  - Tivoli Storage Manager
  - CA ARCserve
  - Legato NetWorker
- IBM Magstar MP 3575
  - Tivoli Storage Manager
  - CA ARCserve
- IBM Magstar 3590
  - Tivoli Storage Manager
  - CA ARCserve

3.6.2 SAN services

Your decision to implement a SAN solution (or not) will usually not be based on technology, but rather on the concrete benefits that a SAN can bring to the enterprise.

Some situations leading to a SAN decision are:

- Need for a specific solution that the present systems cannot fulfill, for example, tape sharing
- Installation of a new machine room, where you want to take advantage of SAN in terms of cabling and distance flexibility
- Implementation of a disaster recovery site, taking advantage of the increased distances and extended connectivity offered by a SAN

In all cases, IBM services will be able to help you in these various phases:

- IBM consulting services for feasibility studies and systems design
- IBM PSS for configuration selections
- IBM PSS for cabling configurations and installation
- IBM PSS for implementing the solutions

Chapter 4. IBM SAN Data Gateway implementation

In this chapter, we describe how we actually used the SAN Data Gateway with tape devices in a lab environment. We had access to a configuration with a Windows NT host, and to another configuration with AIX hosts.

4.1 Device mapping considerations

At the beginning of the implementation you should decide upon the device configuration and mapping. The following description should help you to understand the mapping of SCSI devices to FC targets and Logical Unit Numbers (LUNs).

One of the functions of the SAN Data Gateway is to connect hosts with serial interfaces, such as Fibre Channel, to devices with different interfaces, such as SCSI. In a typical configuration, each Gateway's Storage Area Network interface is connected to either a single host or multiple hosts with the use of Fibre Channel hubs or switches.

It is important to understand how devices attached to the SCSI channels appear to the SAN Data Gateway and to the host systems. This differs from targets that are attached directly to a host system without any intervening gateways, routers or bridges.

4.1.1 Device mapping in FC environments

In SCSI terminology, a tape drive or a disk drive is attached to a "bus" and has a unique address on that bus. There are three parts of the address in a conventional SCSI:

- Bus (or channel)
- Target ID
- LUN

A simple case of two tape drives attached to a single bus is shown in Figure 35.

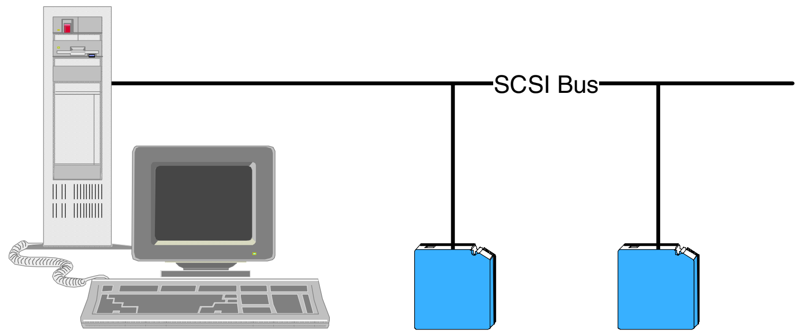

Figure 35. Basic SCSI connection to a system

This configuration has a device map (from the host point of view) that looks as shown in Table 3. Some tape library devices like the Magstar 3570 or 3575 support LUNs for the devices; the device is addressed as LUN 0, the autochanger is available as LUN 1.

Table 3. Target ID and device mapping — native SCSI

Target ID	LUN	Device
0	0	Tape (Target 0)
1	0	Tape (Target 1)
7 (usually 7)		Host system

If a product such as a SAN Data Gateway is placed between the system and target devices, the addressing has another layer. This is because the targets (the tape drives) are not directly attached to the host. Instead, the actual targets are connected to a SCSI adapter that is part of the gateway, and they are not directly attached to the system. Figure 36 shows the device mapping with the additional layer due to the SAN Data Gateway.

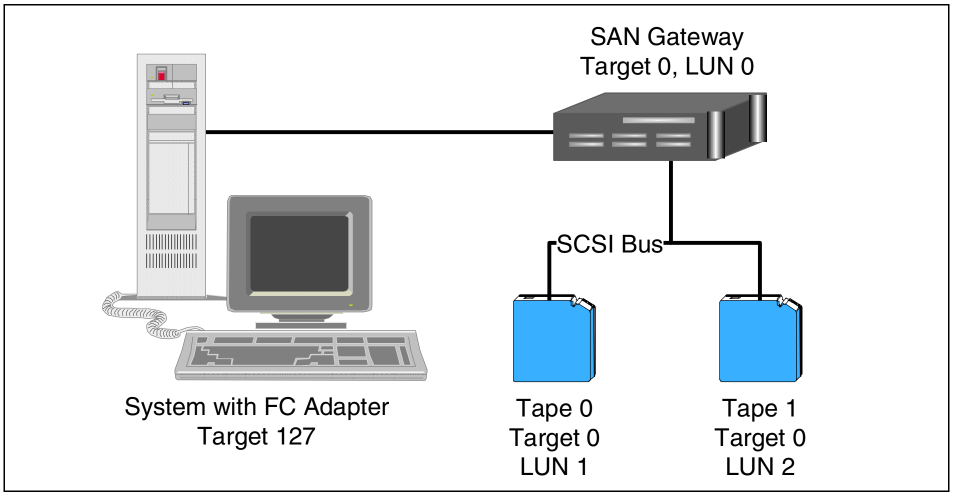

Figure 36. SAN Data Gateway attached through Fibre Channel — host view

The system now only has a single target ID (target 0 in this case) directly attached to the SAN Data Gateway. The gateway forwards commands to and from its targets, the tape drives. However, there is a need to map the devices (tapes) so the host system can use them. This is achieved by using another layer of mapping, LUNs. The device map might now look like Table 4 from the host point of view.

Table 4. Device Map from the Host Point of View — with SAN Data Gateway

Target ID	LUN	Device
0	0	SAN Data Gateway
0	1	Tape 0
0	2	Tape 1
127		Host system

Note that LUN 0 points to the SAN Data Gateway. This allows you to send commands to control the gateway. This is referred to as the Command and Control Interface.

We simplified this by assuming hard IDs for the Fibre Channel devices. A hard ID is a number from 0 to 127 used in place of the World-Wide Name (WWN) assigned to each Fibre Channel device. It requires that SCSI devices are mapped exactly the same way every time the host is booted again. The SAN Data Gateway has the ability to deal with these issues, because it preserves the mapping of device addresses.

When using Fibre Channel hubs or switches, hard IDs are preferred because the host won't lose track of the devices. Scrambled IDs are not a concern when the SAN Data Gateway is used in a Fibre Channel point-to-point configuration.

4.1.2 Mapping devices to targets and LUNs

The host operating system imposes limitations on the number of LUNs and target IDs. The most common SCSI implementation in use (SCSI-2) allows eight LUNs per target ID, so if you want to connect more than seven devices (one LUN is reserved for the gateway), this address limitation needs to be circumvented. Fibre Channel and other serial interfaces implement a newer SCSI specification (SCSI-3), which allows for many more LUNs per target.

The number of LUNs is limited by the host operating system. In case of SCSI-2 protocol, another abstraction is required from the host adapter. This varies with the implementation, but a common method is to create additional "pseudo" target IDs when you use up the eight LUNs. When adding more drives, the host system "sees" target 0 has LUNs 0 through 7, then target 1, LUNs 0 through 7 for the next eight devices, continuing until all devices are accounted for. With a SCSI-2 type mapping, there are 16 target IDs available per bus, but one must be reserved for the host adapter. That allows for 15 targets with 8 LUNs each or 120 total devices.

Different operating systems have different limits on the number of LUNs supported. For example, Windows NT 4.0 had an eight LUN limit until the release of Service Pack 4. The limit is now 255 LUNs per target ID. The increase implies that you need not use "pseudo" targets. When the newer SCSI-3 specifications are implemented by the operating system vendors, it is possible to go beyond the SCSI-2 limits of device addressing. The SAN Data Gateway fully supports SCSI-3 protocol for disk and tape.

Figure 37 explains the mapping of devices to targets and LUNs with an example of a Tivoli Storage Manager for NT installation with IBM Magstar 3570-C12 Tape Library and the SAN Data Gateway. In the upper part you see the Service Point Terminal command output for **fcShowDevs**, and in the lower part you see the Tivoli Storage Manager command output for **q drive.**

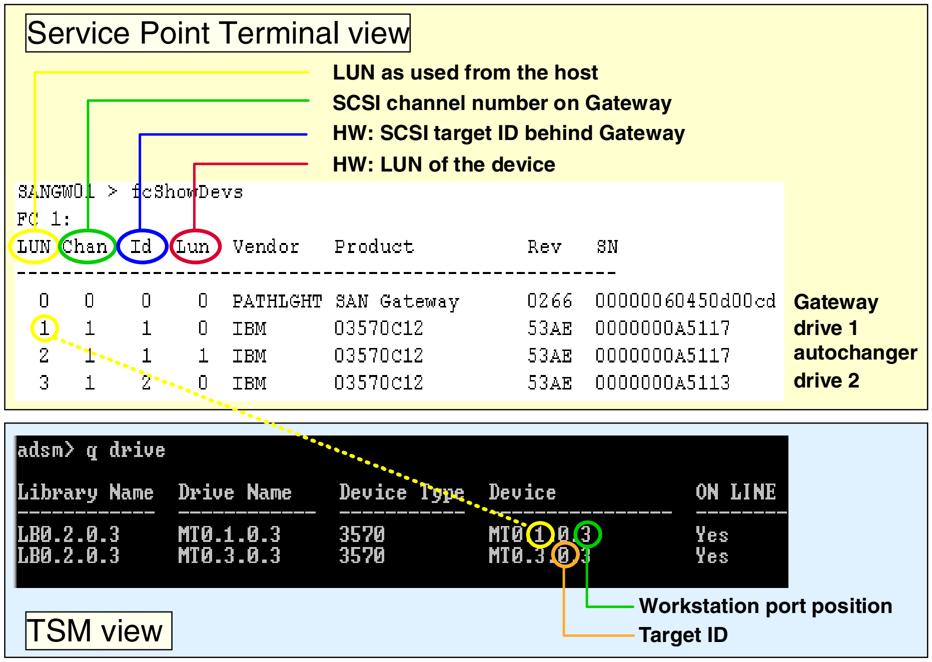

Figure 37. Device mapping views

4.2 Windows NT

This section covers installation of the IBM SAN Data Gateway, attached SCSI devices, and the host adapter in a Windows NT environment.

4.2.1 Setup

The *SAN Data Gateway Installation and User's Guide* contains very good checklists which you should follow during the installation. Instructions on how to install the product are divided into three sections:

- Pre-installation checklist

  Before the IBM Service Representative installs the SAN Data Gateway, check the items on the pre-installation checklist.

- Installation checklist

  The installation checklist has instructions on how to install the SAN Data Gateway for the IBM service representative. For the installation checklist, see the *IBM SAN Data Gateway Service Guide*.

- Post-installation checklist

  After the IBM Service Representative installs the SAN Data Gateway, go through the post-installation checklist.

Test Setup

Figure 38 shows the components and installation setup we used for the SAN Data Gateway installation.

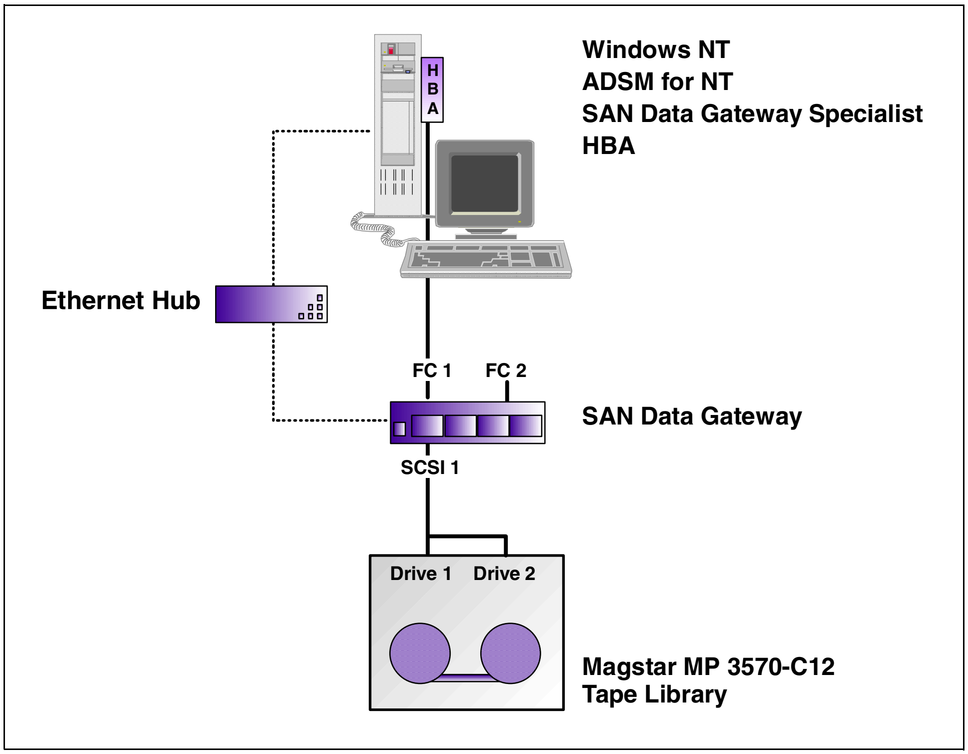

Figure 38. Installation setup

Pre-installation tasks

1. Inform the IBM service representative whether the gateway should be desk top or rack mount installed.
2. Check the Operating System requirements. For Windows NT it is Version 4.0, Service Pack 4 or later, plus required hotfixes. Those hotfixes and a current list of supported platforms, as well as device driver and firmware requirements, are mentioned on the Web site:

   `http://www.ibm.com/storage/SANGateway/`

You can check the level of the OS at boot time, for example, up to the Service Pack level.

3. Install the FC Host Bus Adapter, for Windows NT this is the QLogic card QLA2100.

   During the installation you have to update the QLogic card BIOS, although the readme file of the QLogic installation files does not state this clearly:

   1. Copy the files from the CD-ROM x:\NT\2100 directory to a floppy diskette from which you can boot. If you downloaded a later version of the driver package, extract that to the floppy instead.
   2. Shut down the system.
   3. Install the HBA in an appropriate 32 or 64 bit PCI slot.
   4. Boot the system with the floppy diskette from which you can boot.
   5. From the DOS prompt, enter the command: **ql2xutil /L /F**
   6. When the procedure completes, remove the floppy.
   7. Boot Windows NT.
   8. Extract the file you downloaded (if necessary) to a local, temporary directory, but only if you have downloaded updated device drivers.
   9. From the Windows Start menu, select Settings, Control Panel, then open the **SCSI Adapters** applet.
   10. Click on the Drivers tab. If the driver is already present for the QLA2100, click on it, then press the Remove button.
   11. Click on the Add button.
   12. Click on the Have Disk button. If you have downloaded updated drivers, provide the path to the temporary directory where they are located. If you use the CD-ROM, the drivers are located in the directory x:\NT\ 2100, where x is the letter for the CD-ROM drive.
   13. Confirm the path and wait for the driver to install.
   14. When you are prompted to reboot the system, select “Yes“.

   Verify that the device driver has started. It should show "started" without having any FC cable, gateway, or devices installed yet. You may see a second device entry for the QLogic card stating a Pseudo device. This is correct, and it is there for hotplugging purposes.

4. Check that all host FC cables are either ordered or already installed and mark them with a host server identifier and the SAN Data Gateway identifier.

5. Ensure that the tape system installation has been completed. Install the newest device driver. Mark both ends of each SCSI cable with
   - Target ID and channel number
   - SAN Data Gateway ID and channel number.
6. Decide the SAN Data Gateway Ethernet port configuration parameters:
   - Static IP address
   - Netmask (if required)
   - SAN Data Gateway name (optional, default name is Gateway)

   If the gateway is not on the same TCP/IP subnet as the StorWatch SAN Data Gateway Specialist, assign the default network gateway address and/or route table entries.
7. Run the Ethernet cable from the NT server with the StorWatch SAN Data Gateway Specialist to the Ethernet network hub.
8. Run the Ethernet cable from the network hub to where the SAN Data Gateway will be installed. Leave it unconnected for the IBM service representative to install.

Installation tasks

The IBM service representative handles the installation of the SAN Data Gateway according to the installation checklist in the Service Guide. They will ask you afterwards to verify that all target devices are available to the host systems.

Post-installation tasks

1. Install the StorWatch SAN Data Gateway Specialist software. For the latest version, see the Web site:

   `http://www.ibm.com/storage/SANGateway/`
2. Ensure that the StorWatch SAN Data Gateway Specialist has access to the same Ethernet subnet as the SAN Data Gateway.
3. Start-up the StorWatch SAN Data Gateway Specialist Server.
4. Start-up the StorWatch SAN Data Gateway Specialist Client.

- Connect to the Specialist Server using its name or IP address (if it does not find the Server running on the same system).
- Log on and add a new administrator account and password. This new administrator account will automatically deactivate the default **StorWatch** account which makes sense because of security reasons.
- Log off and log on using the new administrator.

5. From the Client, connect to the desired SAN Data Gateway (use **Connect SAN Data Gateway** from the tools menu group and enter the SAN Data Gateway name or IP address) or discover all installed gateways with function **Discover Net**.

   Figure 39 shows the Specialist after connecting the SAN Data Gateway. The bar at the bottom shows the actual status of the discover net function.

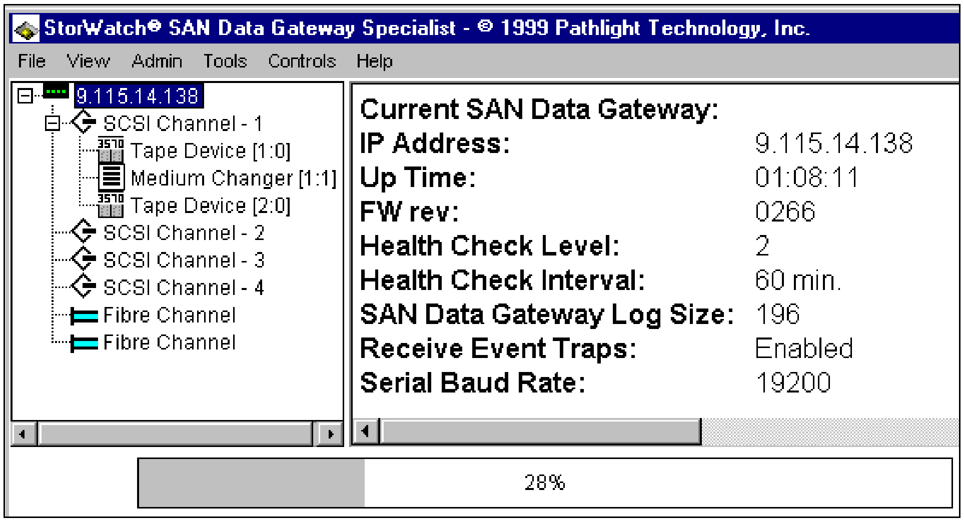

Figure 39. SAN Data Gateway in the StorWatch SAN Data Gateway Specialist

6. Save your user / administrator setup with **Save Setup as...** from the file menu group. This function saves the currently viewed selection of gateways, not the configuration itself.

7. If later firmware for the gateway is available, update each gateway with function **Update Firmware** and **Restart Gateway** from the Control menu group.

   Figure 40 shows the Update Firmware and Restart Gateway functions in the Controls menu group. Figure 41 shows a warning message that comes up after Restart Gateway, since the restart function re-boots the currently selected gateway. This immediately stops all I/O activity.

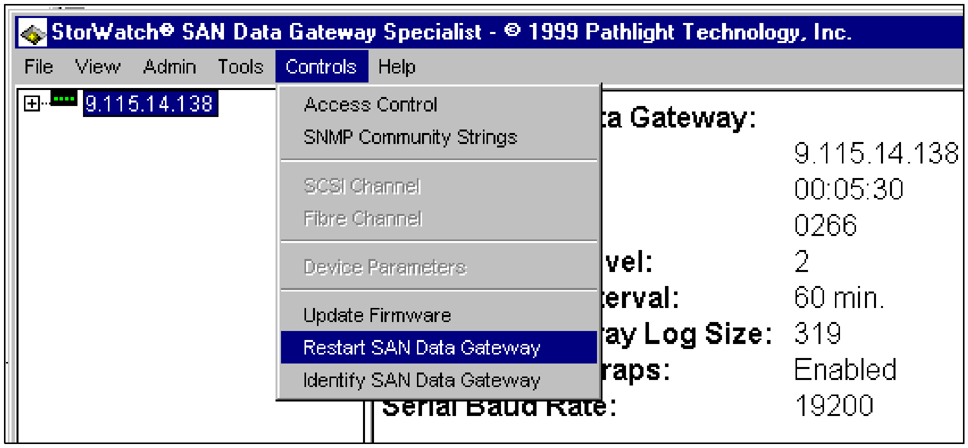

Figure 40. SAN Data Gateway Specialist — update firmware and restart

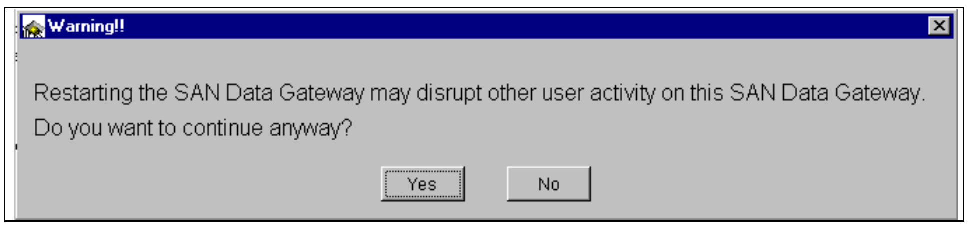

Figure 41. Restart warning message

8. If you have more than one host connected to the SAN Data Gateway, you may need to enable Access Control. Use function **Access Control** in the control menu group to place check marks in the matrix to allow access between the FC connections and the SCSI channels. All combinations are possible.
9. At this time you MUST save your SAN Data Gateway configuration with the function **Save SAN Data Gateway Configuration** in the tools menu group.
10. Become familiar with using the StorWatch SAN Data Gateway Specialist.
11. Use the StorWatch Specialist to monitor and maintain your SAN. There are functions available like event logging, health check, and heartbeat notification.

Save your SAN Data Gateway configuration periodically. You must save your configuration if you make changes to access control or operating parameter settings.

4.2.2 Host considerations

The connection of devices in a SAN changes the way they are seen in the Operating System. This is achieved by using another layer of mapping. The device map has been added to reflect devices in the host which have previously been attached to multiple SCSI cards and are now attached to a single FC HBA.

The address mapping is described in 4.1, “Device mapping considerations” on page 81. Here we want to show how Windows NT defines the devices as shown in Figure 42, and how the StorWatch SAN Data Gateway defines the device shown in Figure 43 on page 92.

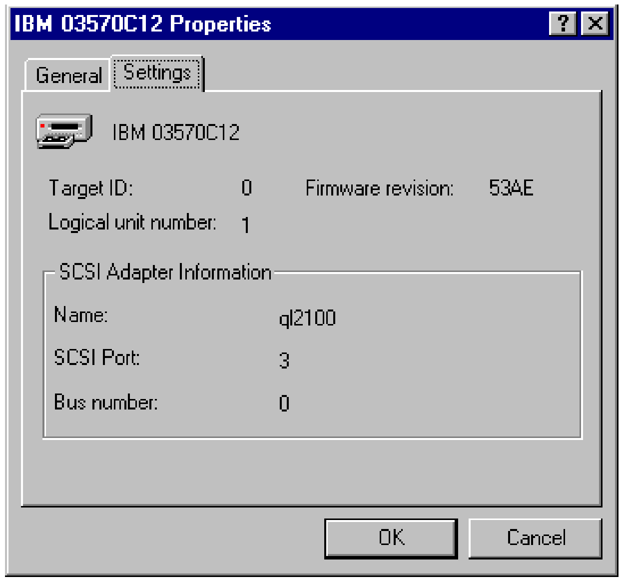

Figure 42. Windows NT SCSI device settings

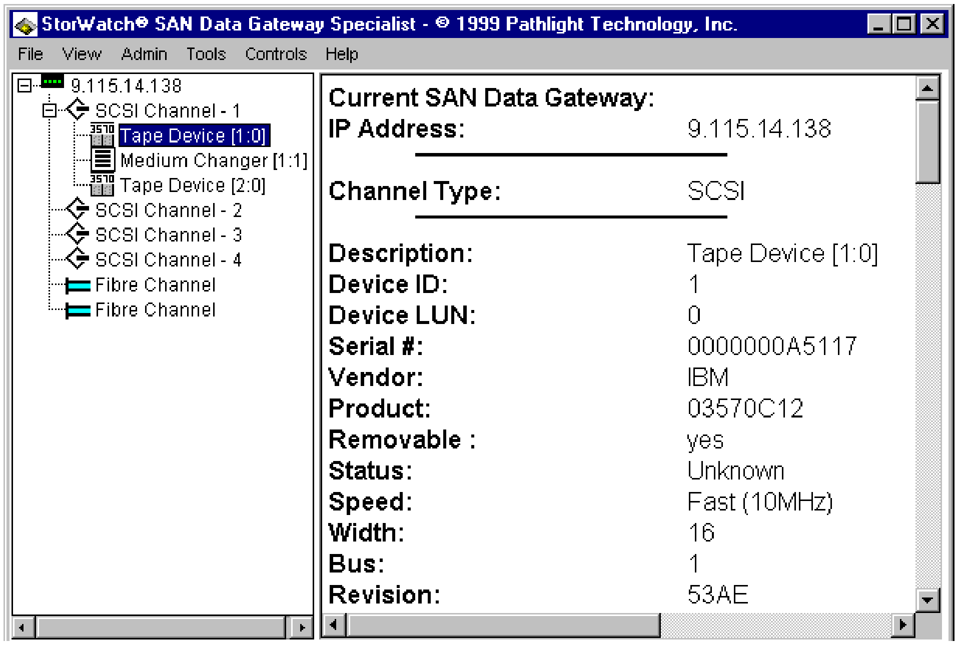

Figure 43. StorWatch SAN Data Gateway — device view

4.3 AIX

We describe here our lab setup and experiences using the SAN Data Gateway with AIX.

4.3.1 Setup

The basic steps that you will need to perform to install the IBM SAN Data Gateway for use under AIX are outlined here:

1. Install the FC Adapter driver on AIX; it is called `fchan0`. Use `smitty` to update the `FCP operation` parameter to `Yes`.
2. Install the FC Adapter card on your RS/6000.
3. Update the FC Adapter card BIOS (or check your BIOS is up-to-date).
4. Connect your SCSI devices to the selected SCSI Port(s) of the SAN Data Gateway.
5. Connect the Fibre Channel cable from the host FC Adapter to the selected SAN Data Gateway fibre port.

6. Power-on, in this sequence:
   1. The SCSI Device(s)
   2. The SAN Data Gateway
7. Re-boot AIX or run `cfgmgr`.

At the end of these steps, AIX should see the SCSI devices on the other side of the SAN Data Gateway as “Available”. The Fibre Channel interface will appear to AIX as `fcnetn`, where n is the network device driver interface number.

Note that steps 2 to 5 are normally done by IBM Service personnel.

8. Set up Ethernet connection for use of the SAN Gateway Specialist.
9. Set up serial the connection for the Service Port interface.

Please refer to the installation checklists in the *Installation and User’s Guide 2108 Model G07* manual for detailed instructions.

Test setup

In our test installation, we had three AIX hosts, connected to one SAN Data Gateway box, with a number of 3570 and 3590 tape drives connected to four SCSI channels on the SAN Data Gateway, as represented in Figure 44 on page 94.

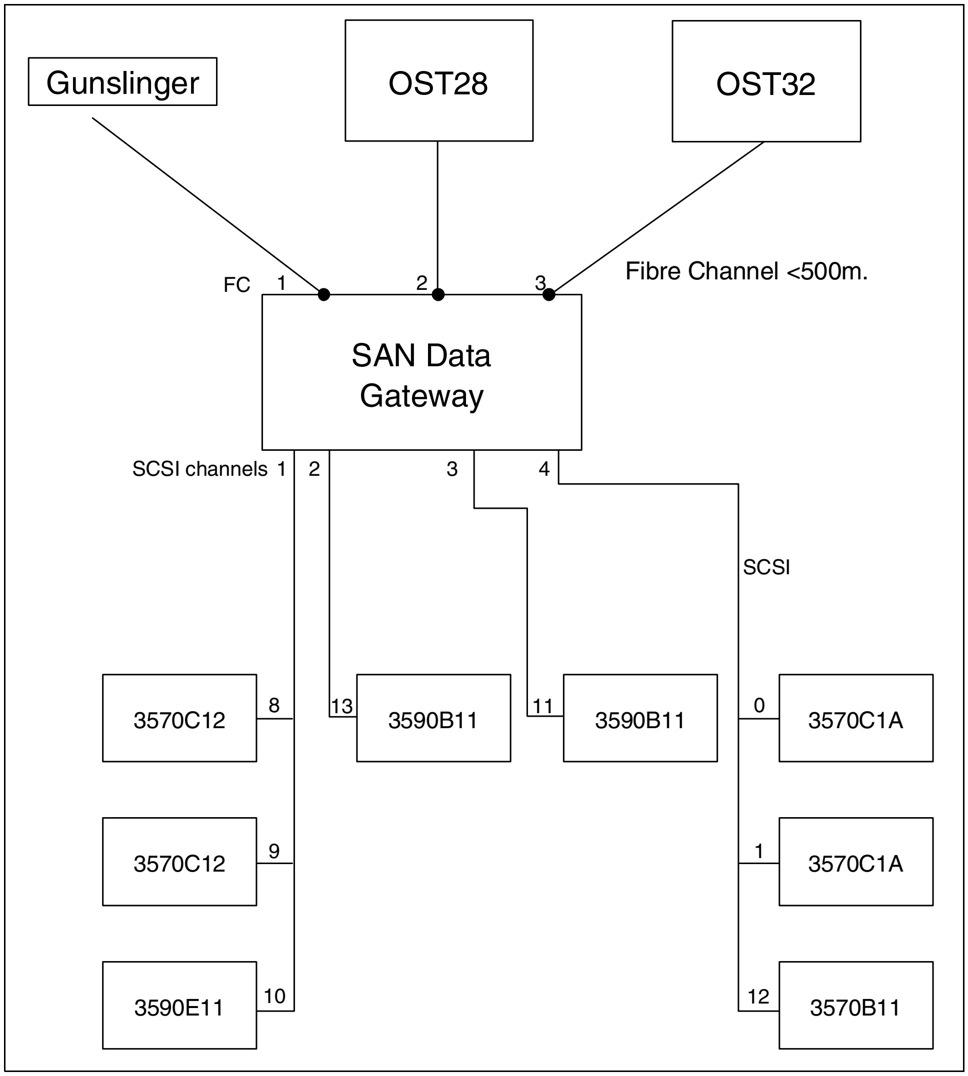

Figure 44. AIX SAN lab configuration

The tape device driver is Atape at the latest level, which include the support for Fibre Channel attached tape drives.

The Access Control were defined so that the ost32 host, connected FC port 1, had access to SCSI ports 3 and 4, while the ost28 host on FC port 2, had access to SCSI channel 1, and the gunslinger host had access to SCSI channel 2. That is, we did SCSI Port zoning so that no devices were seen by two hosts.

The Service Port command **fcShowDevs** thus gives the output shown in Figure 45.

```
266.7 > fcShowDevs
FC 1:
LUN Chan  Id  Lun  Vendor   Product          Rev   SN
----------------------------------------------------
  0   0    0    0  PATHLGHT SAN Gateway      0266  00000060450d00fa
 34   3   11    0  IBM      03590B11         B639  000000039191
 35   3   11    1  IBM      03590B11         B639  000000039191
 36   4    0    0  IBM      03570C1A         5324  000000000344
 37   4    0    1  IBM      03575L18         2.17             0
 38   4    1    0  IBM      03570C1A         5324  000000000346
 40   4   12    0  IBM      03590E11         C523  0000000E0095
 41   4   12    1  IBM      03590E11         C523  0000000E0095
FC 2:
LUN Chan  Id  Lun  Vendor   Product          Rev   SN
----------------------------------------------------
  0   0    0    0  PATHLGHT SAN Gateway      0266  00000060450d00fa
 26   1    8    0  IBM      03570C12         53AE  000000004027
 27   1    8    1  IBM      03570C12         53AE  000000004027
 28   1   10    0  IBM      03590E11         C502  0000000E0086
 29   1   10    1  IBM      03590E11         C502  0000000E0086
 32   1    9    0  IBM      03570C12         53AE  000001310983
 33   1    9    1  IBM      03570C12         53AE  000001310983
FC 3:
LUN Chan  Id  Lun  Vendor   Product          Rev   SN
----------------------------------------------------
  0   0    0    0  PATHLGHT SAN Gateway      0266  00000060450d00fa
 16   2   13    0  IBM      03590B11         B639  000123456789
 17   2   13    1  IBM      03590B11         B639  000123456789
value = 64 = 0x40 = '@'
```

Figure 45. AIX lab fcShowDevs output

The columns shown by fcShowDevs are:

LUN: the SCSI LUN of each target device. This LUN number is internal to the SAN Data Gateway, which is itself LUN 0 in this scheme.

Chan: the channel number of the SCSI interface

Id: the (hardware) target SCSI ID of each device.

Lun: the LUN number of the SCSI device.

The other columns are self-explanatory.

Some things are worth noticing here:

- The LUN numbers are just sequential numbers assigned by the SAN Data Gateway to every device, whereas the Lun (notice the capitalization subtle distinction) is the SCSI LUN used by a given, real device. For example, the 3590B11 device has two LUNs: LUN 0 is the drive, LUN 1 is the Media Changer. This is fixed for this type of device, and cannot be changed.
- The SCSI ID for devices are the ones defined in the device; this is normally selected at installation time.

Seen from one of the hosts, in this example by ost28, which happens to be connected on FC 2, the SMITTY panel displaying tape devices is shown in Figure 46.

```
Command: OK            stdout: yes           stderr: no

Before command completion, additional instructions may appear below.

rmt0 Available 20-60-00-2,0  5.0 GB 8mm Tape Drive
rmt1 Available 10-60-01-8,0  IBM Magstar MP Tape Subsystem (FCP)
rmt2 Available 40-60-00-6,0  Differential SCSI 8mm Tape Drive
rmt4 Available 10-60-01-10,0 IBM 3590 Tape Drive and Medium Changer (FCP)
rmt6 Available 10-60-01-9,0  IBM Magstar MP Tape Subsystem (FCP)
```

Figure 46. AIX lab as seen from ost28

Some notes regarding Figure 46:

- `rmt0` and `rmt2` are directly-attached 8mm drives.
- The devices attached to the SAN Data Gateway have an `(FCP)` indication, meaning Fibre Channel Protocol, which is added by the AIX Atape Device Driver.
- The addresses of the FC-attached devices have the same SCSI IDs and LUNs as seen from the SAN Data Gateway. This is done again by the Atape device driver, which gets this information from the gateway, and puts it in the location field used by `lsdev`, as a convenience. If we were not using Atape, but the standard AIX `ost` (other SCSI tape) driver, we would be in a situation similar to the Windows NT lab, and see all tape drives with the same SCSI ID of 0, and the LUN number as returned by the SAN Data Gateway. This also means that two hosts having access to the same physical device would see it with different SCSI IDs, as these are actually the LUNs returned by the SAN Data Gateway.
- The parent adapter for the FCP devices is called a `fchan0` device.

Note: The AIX location code format for SCSI tape devices has a structure of `AA-BB-CC-S,L` where:

- `AA`: A value of 00 for the AA field indicates the controlling adapter card is located in the CPU drawer or system unit, depending on the type of system.
- `BB`: The BB field identifies the I/O bus and slot containing the card. The first digit identifies the I/O bus. It is 0 for the standard I/O bus and 1 for the optional I/O bus. The second digit is the slot on the indicated I/O bus containing the card. A value of 00 for the BB field indicates the standard SCSI controller.

- `CC`: The CC field identifies the card's SCSI bus that the device is attached to. For a card that provides only a single SCSI bus, this field will be set to 00. Otherwise, a value of 00 indicates a device attached to the card's internal SCSI bus, and a value of 01 indicates a device attached to the card's external SCSI bus.
- `S,L`: The S,L field identifies the SCSI ID and logical unit number (LUN) of the SCSI device. The S value indicates the SCSI ID and the L value indicates the LUN.

For example, the entry labelled `IBM Magstar MP Tape Subsystem (FCP)` on Figure 46 on page 96 means:

- 10-60 is the Adapter location, identifying the drawer and the I/O bus and slot of the FIbre Channel adapter in the RS/6000.
- 01 says that this is an External bus
- 8,0 means SCSI ID 8 and LUN 0 for this device.

4.3.2 Host considerations

Connecting devices across a SAN completely changes the way devices are configured and managed on a host system.

If you don't have a SAN, but normal direct SCSI connections instead, the devices are generally dedicated to one machine, and all the I/O devices connected belong to that host.

In a SAN, a host system will see all the devices connected to its FC Adapter, which appears to the host as one SCSI bus; this can include devices belonging to other hosts.

For example, under AIX, a `cfgmgr` will detect all the devices on all its SCSI busses, and attempt to configure them. And as `cfgmgr` also runs at boot time, there is no way to prevent AIX from doing that.

In a SAN environment, if nothing is done to control this situation, this could lead to:

- Devices belonging to one host being accessed by other hosts, with potentially destructive consequences.
- Devices that were once viewed in a given sequence having that sequence changed just because a new device was added somewhere on the SAN.
- Devices can be known under one name and address from one system, and have another name when viewed from another system, thus leading to confusion and prompting operational errors.

Clearly, this has to be avoided.

There are two aspects to solve to avoid these problems:

- Addressing
- Access control

The SAN Data Gateway addresses these two aspects by:

- Persistent address mapping: The SAN Data Gateway maintains a table of the devices LUNs, so that the hosts always see a given device with the same SCSI address triplet. If a new device is added, it will be assigned a new, unique, LUN. And if a device is removed, its LUN is reserved just in case you reconnect it, and also to avoid that a new device gets the same LUN number.
- Access control (also called zoning): This function of the SAN Data Gateway allows you to assign a SCSI channel to a given FC channel, that is, host. The devices on this SCSI channels are then invisible from the other hosts, thus protecting you from potentially destructive multiple host access. Even though you could let all the hosts see all the SCSI channels, in a normal production mode, you would have only one host accessing a SCSI bus.

  This mode of operation is referred to as port zoning, because you apply restrictions (or authorizations) at the SCSI port, that is, SCSI bus, level. This means that if you have, for example, two tape drives on a SCSI bus, you cannot assign one drive to one host and the other to another host using port zoning.

  There is another type of zoning, called LUN zoning, where the access controls can be applied at the LUN level. This would allow you to have two hosts accessing one drive each on the SCSI bus with two tape drives, for example.

  The SAN Data Gateway does not presently offer LUN zoning, even though the Pathlight Web site mentions that the SAN Data Gateway can do it.

4.4 StorWatch SAN Data Gateway Specialist

This section contains an overview about the functions of the StorWatch SAN Data Gateway Specialist. The Specialist runs on Windows NT, independent of the server platform, to which the gateway is attached.

The StorWatch SAN Data Gateway Specialist is the recommended interface to configure and maintain the gateway remotely on a Windows NT server. The Specialist is outband connected to the gateway through an Ethernet connection with TCP/IP protocol. The SW is delivered as CD with the gateway. You can either install it from the CD or download it from the following Web site:

`http://www.ibm.com/storage/SANGateway/`

In any case, we suggest checking the Web site for the newest version. The installation steps are described in 4.2.1, “Setup” on page 85 under post-installation tasks. In this section we want to describe the different functions of the Specialist software.

Before a Client can manage any SAN gateways, it must connect to the Server component. If the Server and Client are running on the same system, a connection is automatically started when the Client starts. Otherwise, when the application starts, it will show the connection to the server box and you have to enter the name or IP address. After you connect to the server, log on.

At this point, you can:

- Retrieve a saved setup.
- Perform a "Discovery" of all gateways on a subnet.
- Connect directly to a gateway.

Menus allow the user to view the SAN Data Gateways and the devices connected on them, and perform various actions. From the main screen, six major menu groups are available:

- File
- View
- Admin
- Tools
- Controls
- Help

4.4.1 File menu group

The File menu group has commands to load, manipulate, and save setup files. A setup file is your selection of gateways, you want to see in the Specialist. You can also say it is a view. It does not relate to the actual configuration parameters, which you can (and should) save with a different function.

4.4.2 View menu group

The view can show either the entire collection of gateways making up the setup viewed in a hierarchical tree, or it can show the front panel of a single gateway when checking this option. See Figure 47.

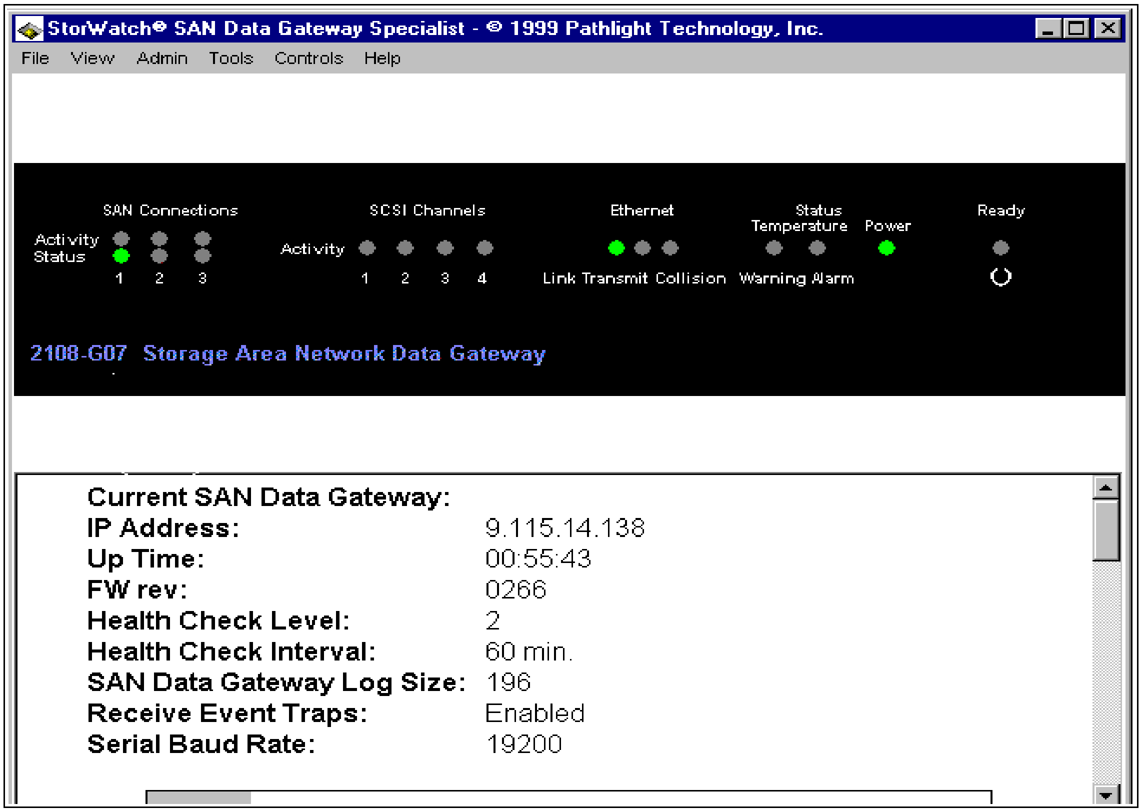

Figure 47. SAN Data Gateway front panel

This view of the panel can be useful to get a quick status of the gateway.

The **Refresh SAN Data Gateway Data** option forces the Server to contact the currently selected gateway and update all the data. It does not restart the gateway. This option is available in the Controls menu group to re-boot the gateway with the current parameters.

4.4.3 Admin menu group

The Admin (administrative) menu group offers control of connecting to the server, logging on and user and password control. Please note that the addition of new users requires administrator privileges.

4.4.4 Tools menu group

The Tools menu group offers options that allow locating and connecting to gateway, and control the reporting of a gateway status.

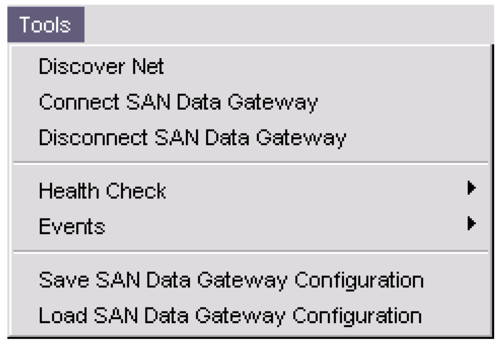

Figure 48. Tools menu group

With the **Discover Net** option in Figure 48 you can discover all gateways in a particular subnet. This option needs some time. An alternative is the **Connect SAN Data Gateway** option that connects to a specific single gateway, which you identify by name or IP address.

The **Health Check** submenus allow you to determine the status of the selected gateway. You can also test target devices and controllers. See Figure 49 for selecting Health Check Intervals and Levels. The Health Check function posts SNMP traps to send status information to management applications. The gateway broadcasts SNMP traps on the local IP network segment.

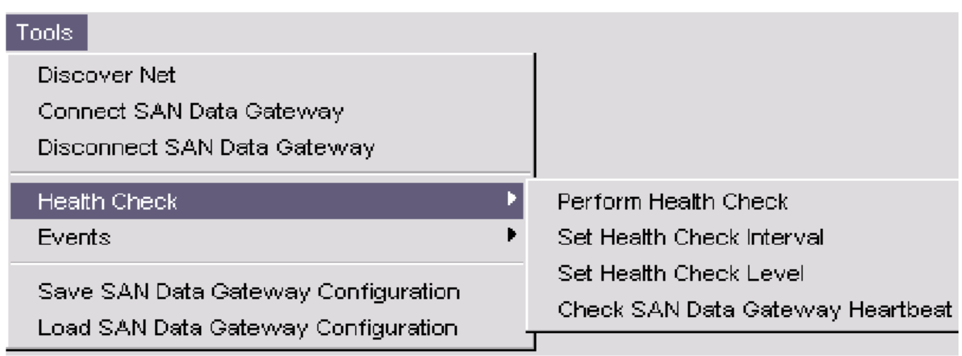

Figure 49. Health check submenu

The **Events** submenu allows you to control how event logs are displayed and how traps are generated.

The **Save SAN Data Gateway Configuration** function copies the vital configuration information from the selected SAN Data Gateway to a file. It is important to preserve the mapping of device addresses. The SAN Data Gateway automatically saves the device map in nonvolatile storage to ensure the persistence between boots and power cycles. To provide a backup of the device maps and other configured parameters, you must **save a configuration** when you:

- Install the SAN Data Gateway
- Add new devices
- Remove devices permanently
- Change Target IDs or LUNs

A configuration, for example, can be reloaded when an IBM Service Representative has to replace the gateway. Save and back up your SAN Data Gateway configurations frequently.

Please note that the function **Save Setup** fulfills a different purpose. The StorWatch SAN Data Gateway Specialist is able to work with multiple gateways which have been previously detected by **Discover Net**. Save Setup saves the currently viewed collection of gateways, so it only saves a view, not the configuration.

The **Load SAN Data Gateway Configuration** function restores a previously saved configuration if you need to replace a gateway. After the load of a configuration, you also have to restart the SAN Data Gateway. A restart will re-boot the selected gateway. This immediately stops all I/O activity.

4.4.5 Controls menu group

The Controls menu group of options allows you to control how the currently selected gateway performs. This includes the definition of the Access Control, the setting of several parameters for SCSI channel, Fibre Channel and target device, and also the Update Firmware and Restart Gateway function.

The **Access Control** function implements the zoning of the SAN configuration based on SCSI and FC ports. This prevents accessing the same target device from more than one host. All combinations are possible.

Figure 50 shows an example where FC 1 has access to SCSI channels 2, 3, and 4, and FC 2 has access to SCSI channels 1 and 4. This means that the devices on SCSI channel 4 can be accessed by multiple hosts, which should be used carefully.

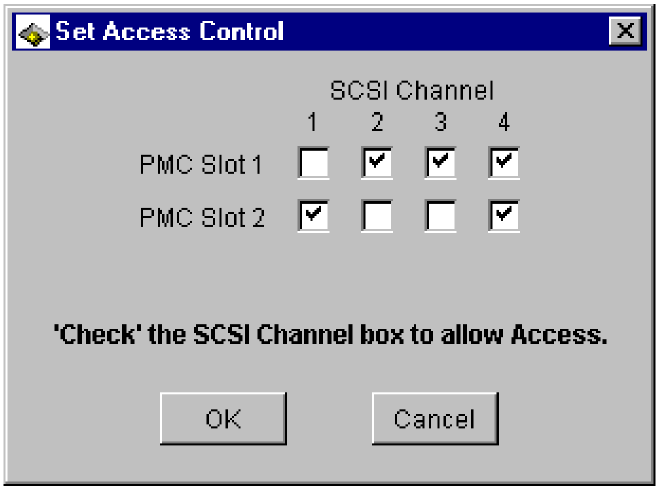

Figure 50. Set access control function

After saving the settings you will get the message shown in Figure 51. The **Restart Gateway** function is also provided in the Control menu group and will immediately stop all I/O activity. Any connected host must stop sending commands to devices attached to the gateway or it will cause time-outs or other failures to be reported on the host system. To reflect the changes in the dialog, you have to refresh the selected SAN Data Gateway.

Figure 51. Warning message after setting new access control

The **Update Firmware** function provides you with a list of firmware files from which you can choose. Select a file and press the **download** button. The selected firmware file is sent to the currently selected device (gateway or target device). You must stop all activity from the connected hosts before you update the firmware for the gateway. After you update the firmware for the gateway, you must re-boot to load the new software.

The **Identify SAN Data Gateway** allows you to visually identify the selected Gateway by a Ready light that flashes rapidly.

4.5 The Service Port interface

The Service Port Interface of the SAN Data Gateway is an RS-232 connection and is used to perform local service and diagnostics. It is primarily intended for IBM Service personnel, but can also be of use if you don’t have access to the SAN Gateway Specialist product.

The connection to the Service Port can be used from a TTY device, or by using a terminal emulator on an NT or AIX system.

Once you are connected to the Service Port, you are in a basic shell interface similar to a UNIX shell, which lets you enter Service Port commands. You can also use the `vi` line editing commands. All these commands are referenced alphabetically in the *User’s Guide 2108 Model G07* manual, Appendix B.

The manual lists all the Service Port commands in alphabetical order. We found it difficult at first to select which command to use for a given function, so we are trying to help you by listing these commands grouped by category, and by telling you which commands to use for the main tasks you will have to accomplish.

Commands grouped by category

The following commands are grouped by category:

- File Management commands:

  Use these commands to manipulate files in the SAN Data Gateway:

  - cd: change directory
  - ls or ll: list the files in the SAN Data Gateway flash memory
  - rm: remove a file

- Connectivity and Control commands

  This set of commands is used for Access Control, to define which host is seeing which SCSI Channel.

  - setFcChanMask: specify which FC channel is allowed to access which SCSI channel
  - setFcFrameSize: change the frame size used on a FC channel. You normally don't have to change this (default is 2048 bytes)
  - setFcHardID: set the Hard AL_PA (Arbitrated Loop Permanent Address) ID for a Fibre Channel, or work with Soft ID method by specifying 255.
  - setFcSplit: assign SCSI channels 1 and 3 for exclusive use by the FC interface 1. SCSI channels 2 and 4 for exclusive use by FC interface 2. This is predefined, simple way to set Access control on the SAN Data Gateway. If you need other assignments, use the setFcChanMask command.
  - setFcNormal: remove the assignments defined with the setFcSplit command.
  - disableCC: make the SAN Data Gateway LUN not available to receive SCSI commands and controls.
  - enableCC: enable the SAN Data Gateway LUN for SCSI commands and controls.

- Device Management commands:

  Use these commands to manipulate the Persistent Address Mapping table in the SAN Data Gateway:

  - mapShowDatabase: show the persistent address mapping table, including all devices, those presently connected as well as the ones which may have been previously connected.

    - mapRebuildDatabase: erase the Persistent Address Map database. A reboot command is then needed to reconstruct a new database in the SAN Data Gateway.
    - mapWinnowDatabase: eliminate inactive device entries from the Persistent Address Map database.
    - mapShowDevs: show a cross reference map of device addresses, ordered by Device ID.
    - scsiRescan: do a SCSI scan on one or all SCSI channels. Use if you add a new SCSI device, or if you have to replace a SCSI device. The rescan will update the configuration data in the gateway; the replacement device will have the same LUN as the old one.
    - scsiShow: list all SCSI channels and all the attached SCSI devices. ID/LUN.
    - targets: lists each target device attached to the gateway, with descriptions as returned by SCSI Inquiry.
    - fcShowDevs: lists all attached target devices with both their SCSI addresses and the Fibre Channel LUN assigned to them.
- Ethernet commands:

  Use these commands to manage Ethernet access to the SAN Data Gateway using the SAN Gateway Specialist. Please refer to the manual for details:

    - ethAddrSet
    - ethEnable
    - ethAddrGet
    - gateAddrSet
    - gateAddrGet
    - ethDisable
    - ethEnable
    - host
    - hostNameSet
    - route
    - trapDestAdd, trapDestRemove, trapDestShow
    - userAdd
    - userList
    - userDelete
    - ping

- Healthcheck commands:

  Use these commands to control the Health checking behavior, like level of checking and interval. The Health Check function broadcasts SNMP traps on the local UP network segment, and also sends them to the SAN Gateway Specialist server. You can specify additional recipient hosts for SNMP events by the trapDestAdd command:

  - hlthChkintervalSet
  - hlthChkLevelGet
  - hlthChkLevelSet
  - hlthChkNow

- Diagnostic and Control commands:

  These commands are normally used by service personnel. The showBox command gives you a graphical view of the SAN Data Gateway box:

  - reboot
  - reset
  - showBox
  - snaVersion
  - sysVpdShow
  - loggerDumpCurrent
  - loggerDump
  - ifShow
  - fcShow
  - csEtimeShow
  - ridTag
  - initializeBox

- OnBoard diagnostic commands:

  These are additional commands, for use by trained service personnel only, which are not documented in the manual.

Task-oriented approach

Tasks and their associated commands are listed in Table 5.

Table 5. Tasks and associated commands

Task	Command	Remarks
Set Gateway name	hostNameSet	Default is Gateway, pick a name according to your standards
Define a new administrator	userAdd	Default is StorWatch, change it for security reasons
Set Gateway Ethernet address and submask	ethAddrSet	See your network administrator
Check Ethernet connectivity	ping	
Display picture of the gateway	showBox	
List target SCSI devices	targets	
List all SCSI channels and the attached SCSI devices	scsiShow	
List all target devices with both SCSI IDs and FC LUN	fcShowDevs	
Rescan all SCSI channels	scsiRescan	Use this after installing new devices
Show a cross-reference of device addresses by Device ID	mapShowDevs	
Display Persistent Address Mapping table	mapShowDatabase	To get SCSI mappings
Define Access Control restrictions	setFcChanMask	Define which host can see which SCSI port

Chapter 5. Scenarios and sample configurations

In this chapter, we present selected sample configurations, using the technologies and functions we have introduced in the previous sections. The samples are selected with a specific customer benefit in mind, like overcoming distance limitations, improving availability, and enhancing connectivity.

5.1 Distance solution

Reasons for using distance solutions could be tape vaulting, centralization of the tape pool, or layout flexibility in the machine room.

The SAN Data Gateway solves problems like the 25 m cable length restriction of SCSI. You can have distances between the gateway and the FC-AL adapter of up to 500 m with multi-mode fiber cables. The attaching fiber cable must be 50.0/125 micrometers for distances of up to 500 meters, or 62.5/125 micrometers for distances up to 175 meters.

For longer distances up to 11 kilometers, you can choose the IBM Fibre Channel Hub. You need two hubs which are connected by the longwave GBIC ports to get the long distance. A short-wave connection of up to 500 m is used between hub and the SAN Data Gateway and the hub to the host.

The IBM Fibre Channel Switch together with the IBM SAN Data Gateway Router are further possibilities to avoid the 25 m cable restriction for SCSI attached devices. The switch supports both shortwave as well as longwave optical connections.

Figure 52 shows an example for an attachment of two Magstar MP 3570 tape libraries through a gateway to a remote Windows NT system. The FC attachment in this case will give you 500 m of distance. To go even further, you need a hub, which provides an additional 10 km of distance.

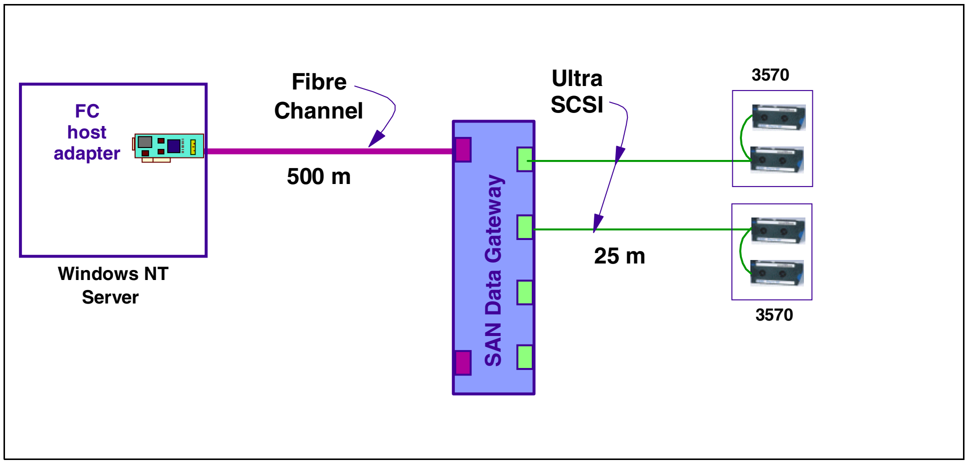

Figure 52. Distance solution with SAN Data Gateway

5.2 Connectivity solution

Connectivity has different aspects. Configurations can get quite complicated when you want to connect multiple hosts to a number of drives. Sometimes it is simply a cabling problem, especially when you want to switch devices between different hosts for example for fail-over purposes. Magstar 3590 helps in this area, since it has two SCSI connectors with which it can be connected to two hosts. Most tape drives only have one SCSI connector per drive.

- **Multiple devices attached to a single (or few) hosts**

  The advantage of SAN in this configuration is that it reduces the number of host adapters you need to attach the devices.

  The SAN Data Gateway currently supports 4 SCSI channels with 3 FC-AL adapters. The number of drives to attach to them is more or less restricted by the effective gateway throughput of around 110 MB/sec.

  The use of hubs in SAN configurations increases the device connectivity even more, but hubs have some issues with multiple hosts on the FC-AL loop such as LIP and arbitration.

  You can also use the Fibre Channel Switch with the SAN Data Gateway Router or a hub for device fan-out, but each router only supports 2 SCSI channels, so the multiplication effect of the router is limited.

- **Multiple hosts attached to a few devices or libraries**

  This sequential device sharing can be required with fail-over configurations where you want to connect a second host to the same tape library, so you can switch to the second host without recabling. Other examples are test machines and application with direct client access to tape.

  The gateway currently supports up to 3 FC-AL ports, which means attachment to three hosts. With either 2 or 3 FC-AL ports, the gateway doesn't really help for host fan-out.

  The IBM Fibre Channel Storage Hub may be a solution for host fan-out, since it has 7 ports, but the issues with multiple hosts on the FC-AL may restrict this possibility.

  The use of the IBM Fibre Channel Switch with either 8 or 16 ports increases the host fan-out which means the number of host connections of SAN configurations. You also can cascade multiple switches to get an even higher number of host connections.

Figure 53 shows a Magstar MP 3575 shared between two hosts. The connection of multiple hosts to a single drive in the 3575 is not possible without a SAN solution, because each drive only has one SCSI connector. The access control function ensures that each host sees only its devices. Access control can be implemented with the SAN Data Gateway Specialist and is also called zoning. See 3.4.3, “Zoning or access control” on page 65 for an explanation of the access control function.

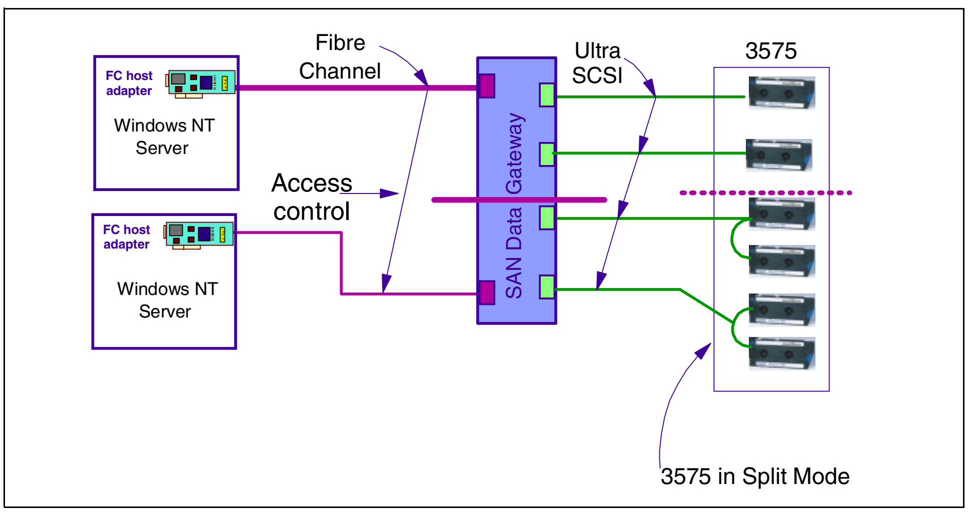

Figure 53. Magstar MP 3575 shared between two hosts

Figure 54 shows an example where a tape library is switched between two hosts, for example, after the second host or one library failed.

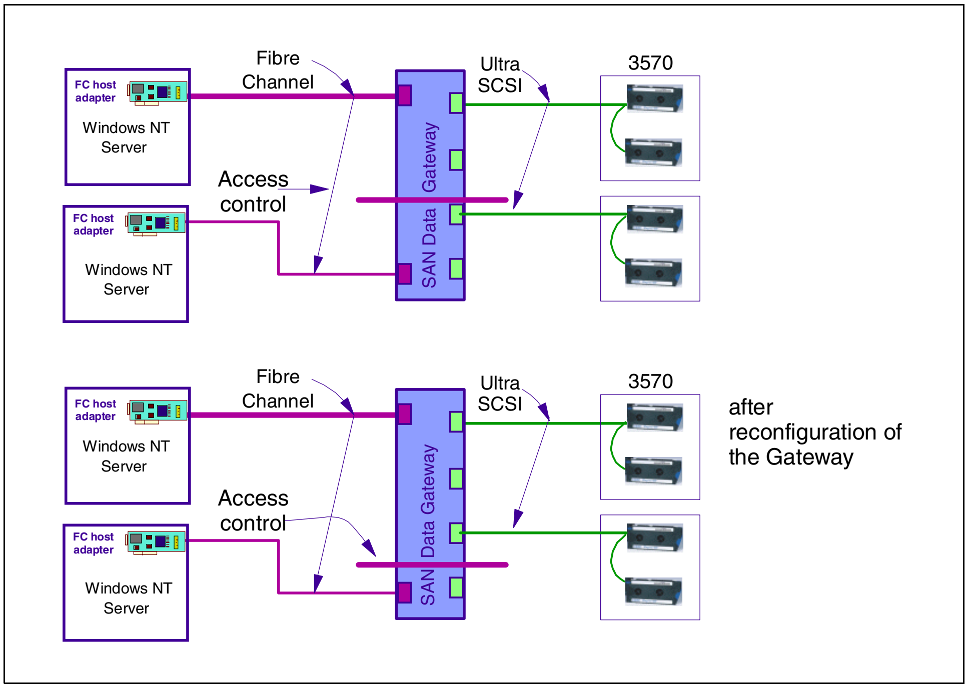

Figure 54. Magstar MP 3570 switched between two hosts

5.3 High availability solution

SAN elements added to a configuration provide you with many advantages, like extended distance and less host adapters, but also introduce additional elements that can fail.

This configuration gives you two access paths to the tape drives from each host. This means that in case of a failure in a Fibre Channel host adapter or in one of the SAN Data Gateways, you can still reach the tape devices from your production host.

This is possible, thanks to the two SCSI ports provided by the 3590 tape drive, thus allowing each drive to be connected to two SAN Data Gateways.

5.3.1 High availability for SAN components

Let us take a simple case where you have only one host system, and configure your SAN boxes without any single point-of-failure in the SAN, as illustrated in Figure 55.

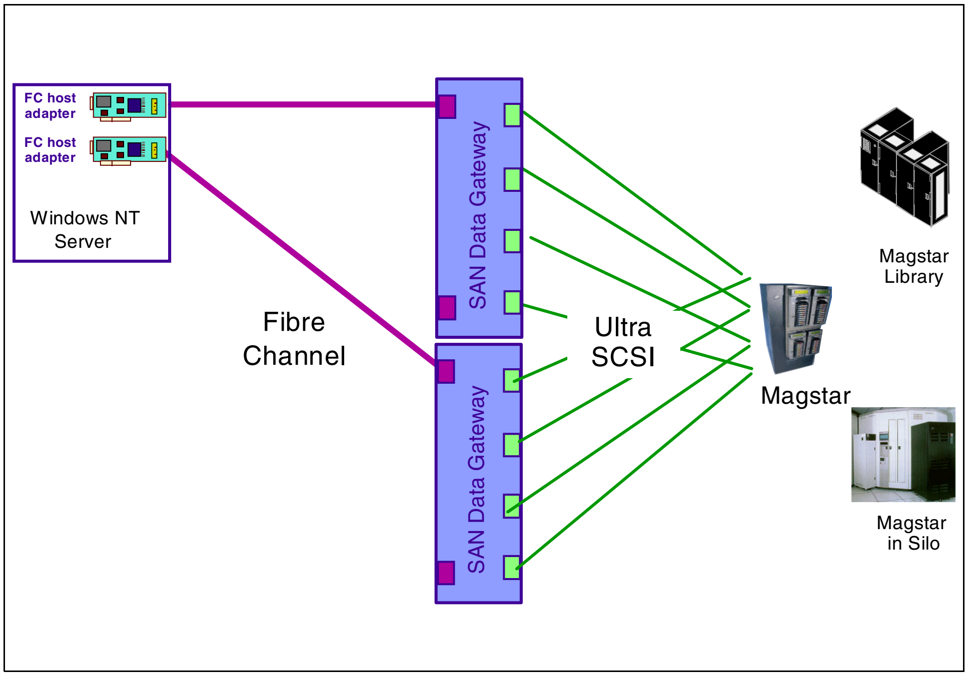

Figure 55. Avoiding single point-of-failure SAN elements

In normal operations, you would have the tape devices visible on both paths, but one given drive would be defined on one path only to the application. In Tivoli Storage Manager terms, for example, you would have one Library definition for each real tape library, and you would define the library drives with the device address corresponding to one of the paths, the “production” path. In case of failure of the production path, you would basically tell Tivoli Storage Manager to use the other path by updating the drives definitions to change the device address to the address going over the other path.

If you do not have two SCSI ports per tape device, that is, any other tape drive, you could have something equivalent using Fibre Channel Switches and the Gateway Router instead of the SAN Data Gateway. See Figure 56.

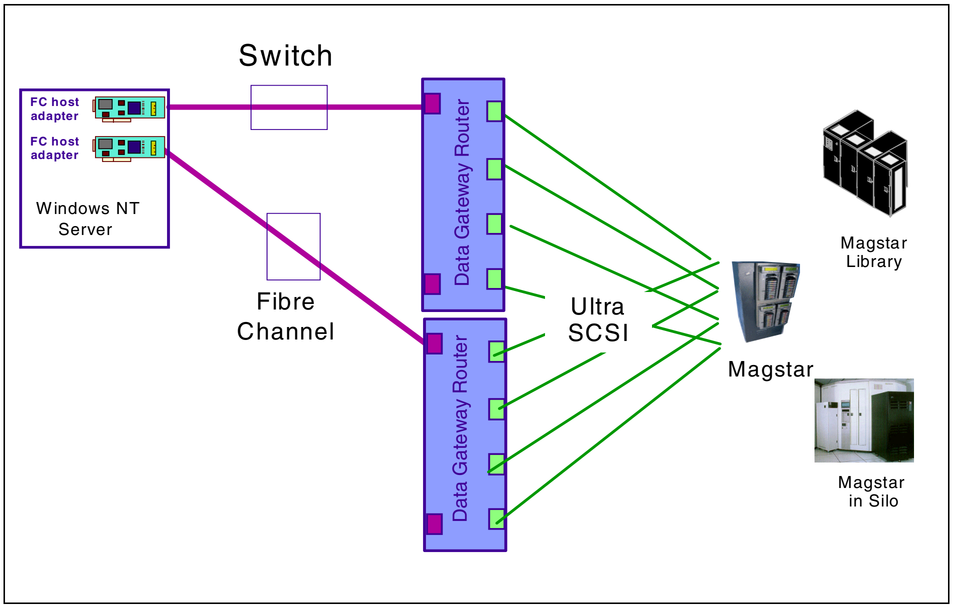

Figure 56. SAN high availability using switches

5.3.2 High availability for hosts in a SAN

If we now consider two hosts systems, you can establish a SAN configuration that gives you the necessary connectivity to implement a high availability solution where one host can take over the applications of the failing host, and you still have the necessary connectivity to your tape drives and libraries from the stand-by host. In this case, you would probably also avoid single points-of-failure in the SAN elements, so we have the same dual pathing as in the previous example, as shown in Figure 57.

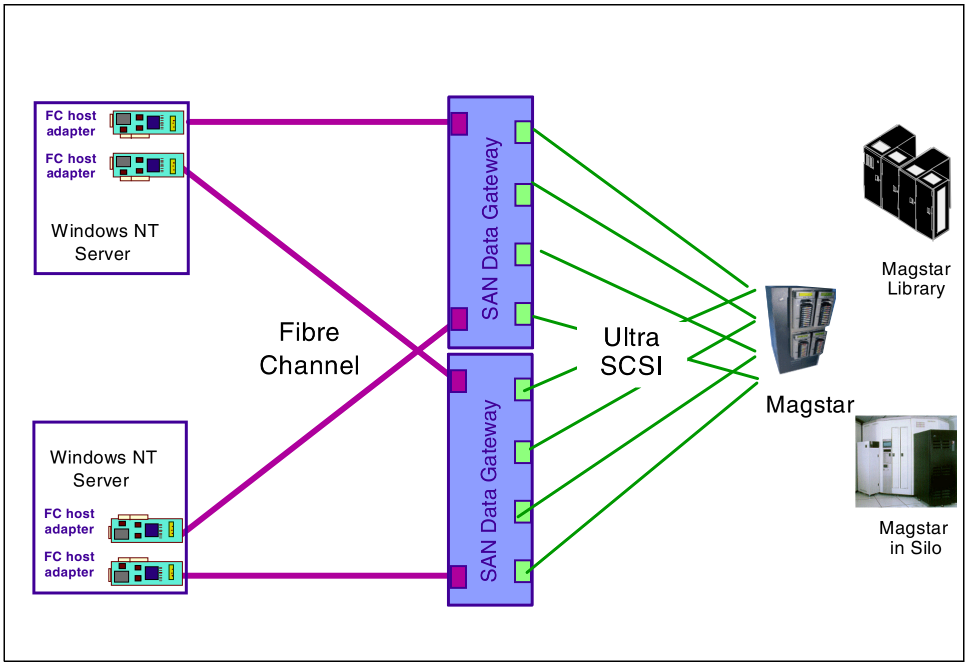

Figure 57. SAN configuration without single points-of-failure

In the case of two hosts, you would have all the tape devices visible from both hosts, but would probably not start the tape application on the stand-by host.

In case of failure of the production host, the operational tasks you would have to perform are:

- Make the tape devices online on the stand-by host
- Bring up the application on the stand-by host

Note that this addresses only the tape connectivity issues. For application take-over, you also need access from the stand-by host, to the disk drives used by the application on the production system. For example, in the case of Tivoli Storage Manager, you would need access to at least the database and recovery log volumes, and preferably also the disk storage pool volumes. This disk connectivity from the two hosts can be achieved either by having the disk devices connected on the SAN, for example, by using an additional SAN Data Gateway, or by another type of multihost connection from the disk units, like SSA disks connected to both hosts.

You could also implement this high availability solution using a SAN Data Gateway Router and a Fibre Channel Switch, as illustrated in Figure 58.

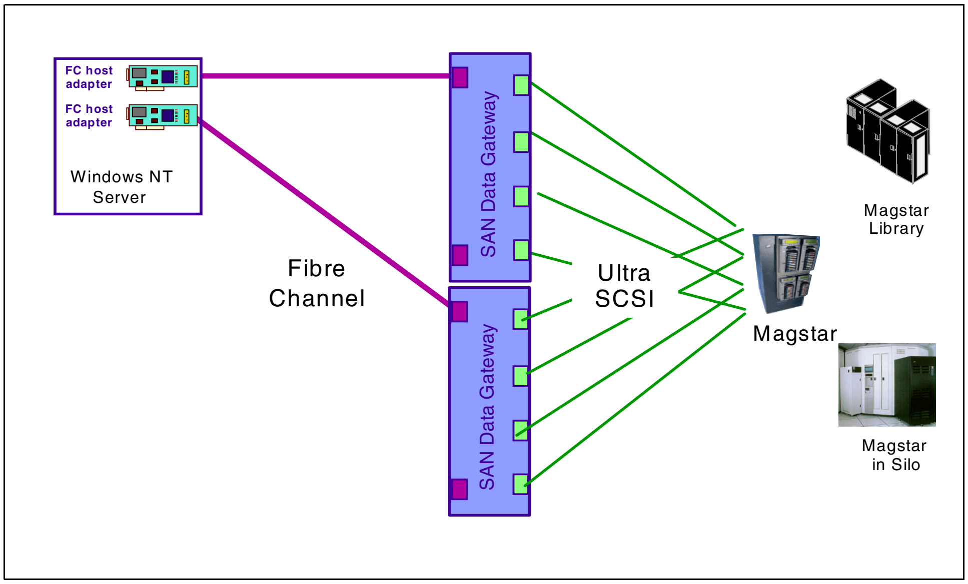

Figure 58. High availability using a Fibre Channel Switch

5.4 Disaster recovery solution

The extended distances of up to 10 kilometers provided by a SAN can be used to provide a disaster recovery solution with a remote site. An example of such a configuration is shown on Figure 59 on page 117, illustrating how a tape library and a host can be installed remotely from the production site.

Extended distance to the recovery site is obtained by using Fibre Channel Hubs. This is necessary because the host Fibre Channel adapters in this example support only short-wave fiber, which has a maximum length of 500 meters, whereas the hub-to-hub links can use long-wave fiber, allowing up to 10 kilometers.

The tape library is placed at the remote site, as well as a backup host system. In case of a disaster at the primary operations site, services and applications can be restarted at the remote site, all the tape data being immediately available.

This example, however, does not protect you against a disaster at the remote site that would destroy all your tape volumes. So you would still need to produce disaster copy volumes, and place them at another site. One solution that would not need any physical movement of cartridges would be to have another tape library at the primary operations site.

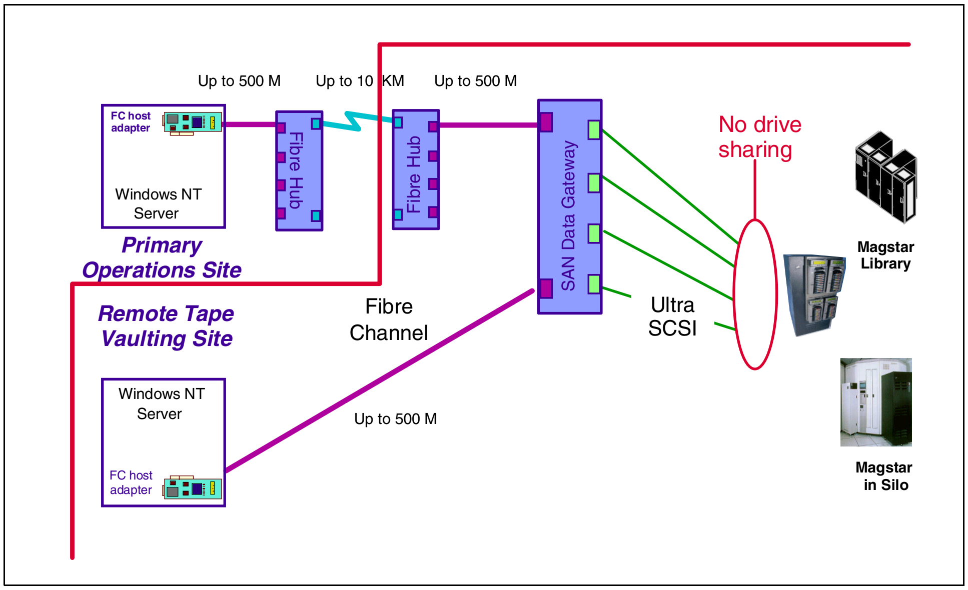

Figure 59. SAN solution for disaster recovery

5.5 Tape sharing solutions

In the next few sections, we outline the solutions available in a SAN environment for tape sharing issues.

5.5.1 Static tape sharing

With static tape sharing, we mean the time based partitioning of either whole tape libraries or tape drives between different hosts. Gateways, hubs, and switches can be used to implement the attachment of the device to multiple hosts, which can allow access to the devices by reconfiguration. It is not necessary to re-cable the devices, only a reconfiguration of the SAN component and, of course, the required actions must be done in the application to make a device available / unavailable.

Figure 54 on page 112 is an example of the sharing of a Magstar 3570 tape library between two Windows NT systems.

5.5.2 Homogeneous tape sharing

Tivoli Storage Manager with its tape sharing function supports heterogeneous library sharing so that multiple Tivoli Storage Manager servers of mixed platforms can share the same library. It allows you to have essentially as many Tivoli Storage Manager server as you need, without needing to install tape libraries on each server machine. Reasons for using multiple additional Tivoli Storage Manager server hosts are, for example, load balancing, testing, and server take-over. An additional advantage is that the Tivoli Storage Manager tape sharing function reduces the numbers of tape drives required in a configuration, since the drives can be shared between the servers.

Figure 60 shows a tape pooling solution with the attachment of seven Tivoli Storage Manager/NT servers attached to a single Magstar 3494 tape library through a Fibre Channel Switch and a SAN Data Gateway. The servers do not have to be of the same platform.

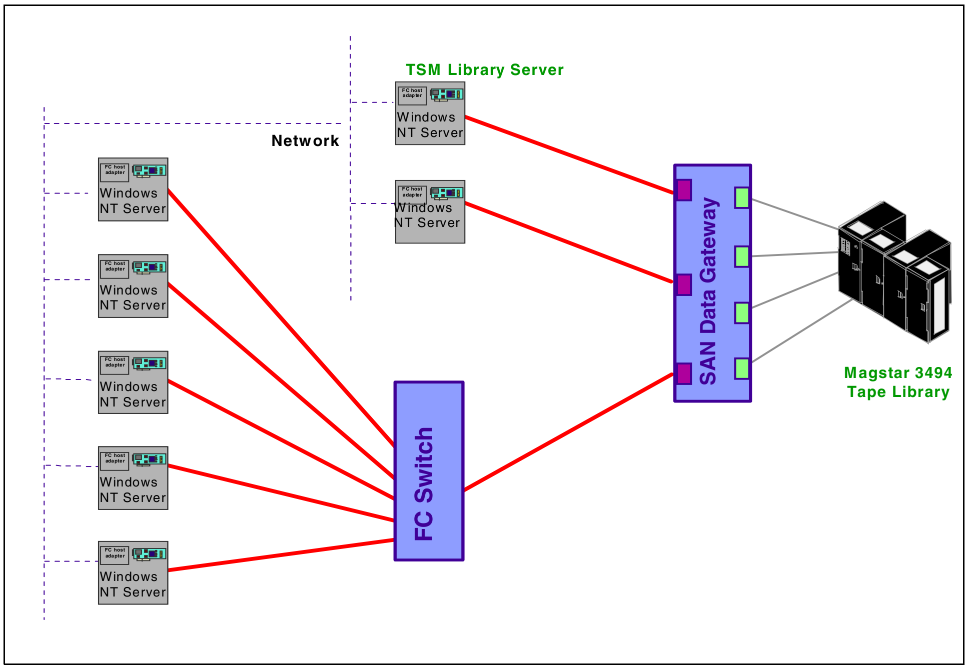

Figure 60. Tape pooling solution

5.5.3 Heterogeneous tape sharing

Heterogeneous Tape Sharing has a Library Server host controlling the tape devices and libraries, allocating tape units, and managing tape cartridges for the clients on the SAN. These clients are in principle any tape application: backup products like Tivoli Storage Manager, Legato, and others, user-developed tape applications (RYO), and also basic operating system backup utilities, like mksysb in UNIX. This is illustrated on Figure 61.

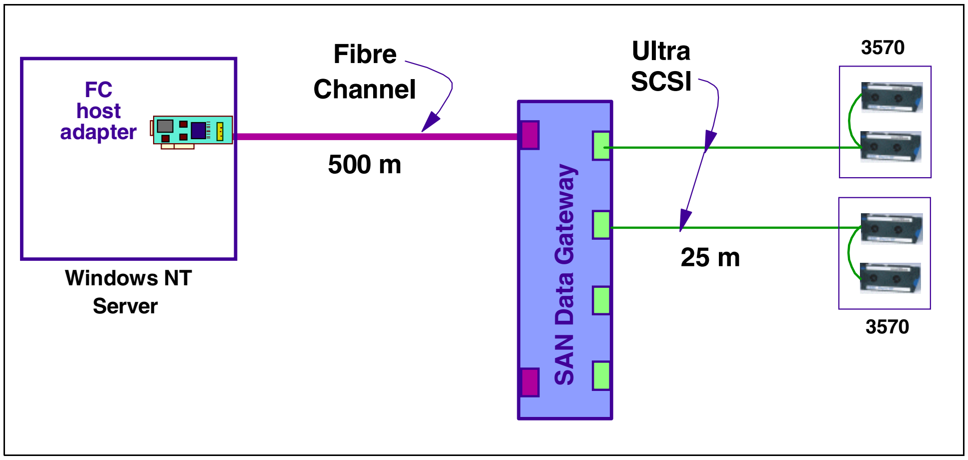

Figure 61. Heterogenous tape sharing

5.6 Disk and tape on SAN

If you have a SAN with SAN-attached disks as well as SAN-attached tape devices, you can use advanced functions like third-party copy between disks and tapes, but you can also only take advantage of the connectivity SAN gives you to implement a generalized fail-over system.

Many installations have fail-over systems for the critical applications. With direct SCSI or SSA connections, where the devices can only be reached from one or two systems, this often forces you to have several of these fail-over systems, one per application. You don't need to do this any more when you have connectivity to all the SAN devices from a given host. One fail-over system can be used for all the applications having their data on the SAN, for example, for Tivoli Storage Manager server and database servers, whether they require tape devices or not.

Appendix A. Fiber Channel discussion

Today, the industry considers Fibre Channel (FC) as the architecture on which most SAN implementations will be built. Fibre Channel is a technology standard that allows data to be transferred from one network node to another at very high speeds. Current implementations transfer data at 100 MB/sec. The 200 MB/sec and 400 MB/sec data rates have already been tested. This standard is backed by a consortium of industry vendors and has been accredited by the American National Standards Institute (ANSI). Many products are now on the market that take advantage of FC's high-speed, high-availability characteristics.

Note that the word *Fibre* in Fibre Channel is spelled in the French way rather than the American way. This is because the interconnections between nodes are not necessarily based on fiber optics, but can also be based on copper cables.

Some people refer to Fibre Channel architecture as the Fibre version of SCSI. Fibre Channel is an architecture used to carry IPI traffic, IP traffic, FICON traffic, FCP (SCSI) traffic, and possibly traffic using other protocols, all at the same level on the standard FC transport. An analogy could be ethernet, where IP, NetBIOS, and SNA are all used simultaneously over a single Ethernet adapter, since these are all protocols with mappings to ethernet. Similarly, there are many protocols mapped onto FC.

FICON is expected to be the standard protocol for S/390, and FCP is the expected standard protocol for the non-S/390 systems, both using Fibre Channel architecture to carry the traffic.

In the following sections, we will introduce some basic Fibre Channel concepts, starting with the physical layer and proceeding to define the services offered.

A.1 Layers

Fibre Channel is structured in independent layers, as are other networking protocols. There are five layers, where 0 is the lowest layer. The physical layers are 0 to 2, the upper layers are 3 and 4.

A.1.1 Lower layers

- **FC-0** defines physical media and transmission rates. These include cables and connectors, drivers, transmitters, and receivers.

- **FC-1** defines encoding schemes. These are used to synchronize data for transmission.
- **FC-2** defines the framing protocol and flow control. This protocol is self-configuring and supports point-to-point, arbitrated loop, and switched topologies.

A.1.2 Upper layers

Fibre Channel is a transport service that moves data fast and reliably between nodes. The two upper layers enhance the functionality of Fibre Channel and provide common implementations for interoperability.

- **FC-3** defines common services for nodes. One defined service under development is multicast, to deliver one transmission to multiple destinations.
- **FC-4** defines upper layer protocol mapping. Protocols such as FCP (SCSI), FICON, and IP can be mapped to the Fibre Channel transport service.

A.2 Topologies

Fibre Channel interconnects nodes using three physical topologies that can have variants. Topologies include:

- ***Point-to-point*** — The point-to-point topology consists of a single connection between two nodes. All the bandwidth is dedicated for these two nodes.
- ***Loop*** — In the loop topology, the bandwidth is shared between all the nodes connected to the loop. The loop can be wired node-to-node; however, if a node fails or is not powered on, the loop is out of operation. This is overcome by using a hub. A hub opens the loop when a new node is connected and closes it when a node disconnects.
- ***Switched*** — A switch allows multiple concurrent connections between nodes. There can be two types of switches, circuit switches and frame switches. Circuit switches establish a dedicated connection between two nodes, whereas frame switches route frames between nodes and establish the connection only when needed. A switch can handle all protocols, as it does not look at the Fibre Channel layer FC-4.

A.3 Classes of Service

Fibre Channel provides a logical system of communication called *Classes of Service*, which are allocated by various Login protocols. The following five classes of service are defined:

- **Class 1** — Acknowledged Connection Service

  Dedicates connections through fabric equivalent of a dedicated physical link and delivers frames with acknowledgment in the same order as transmitted.

- **Class 2** — Acknowledged Connectionless Service

  Multiplexes frames from multiple sources with acknowledgment. Frame order is not guaranteed.

- **Class 3** — Unacknowledged Connectionless Service

  Same as class 2, without frame acknowledgment. Flow has to be controlled at buffer level.

- **Class 4** — Fractional Bandwidth Connection Oriented Service

  Same as class 1, but with only a minimum of bandwidth guaranteed. If sufficient bandwidth is available, class 2 and 3 frames will share connections.

- **Class 6** — Simplex Connection Service

  Same as class 1, but also provides multicast and preemption.

A.4 SAN components

As mentioned earlier, the industry considers Fibre Channel as the architecture on which most SAN implementations will be built, with FICON as the standard protocol for S/390 systems, and FCP as the standard protocol for non-S/390 systems. Based on this perception, the SAN components described in the following sections will be Fibre Channel based.

A.4.1 SAN servers

The server infrastructure is the underlying reason for all SAN solutions. This infrastructure includes a mix of server platforms such as Windows NT, UNIX (various flavors) and OS/390. With initiatives such as Server Consolidation and e-business, the need for SAN will increase. Although the early SAN solutions only supported homogeneous environments, SAN will evolve into a truly heterogeneous environment.

A.4.2 SAN storage

The storage infrastructure is the foundation on which information relies and therefore must support a company's business objectives and business model. In this environment, simply deploying more and faster storage devices is not enough; a new kind of infrastructure is needed, one that provides more enhanced network availability, data accessibility, and system manageability than is provided by today's infrastructure.

The SAN meets this challenge by liberating the storage device, so it is not on a particular server bus, and attaches it directly to the network. In other words, storage is externalized and functionally distributed across the organization. The SAN also enables the centralizing of storage devices and the clustering of servers, which makes for easier and less expensive administration.

A.5 SAN interconnects

The first element that must be considered in any SAN implementation is the connectivity of storage and server components using technologies such as Fibre Channel. The components listed below have typically been used for LAN and WAN implementations. SANs, like LANs, interconnect the storage interfaces together into many network configurations and across long distances.

Much of the terminology used for SAN has its origin in IP network terminology. In some cases, the industry and IBM use different terms that mean the same thing, and in some cases, mean different things. For example, after reading the sections A.5.10, "Bridges" on page 127 and A.5.11, "Gateways" on page 127, you will find that the IBM SAN Data Gateway is really a bridge, not a gateway.

A.5.1 Cables and connectors

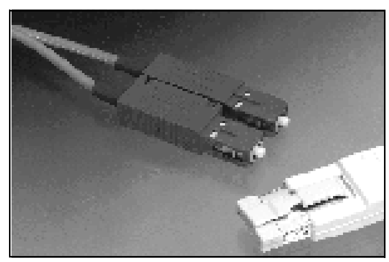

As with parallel SCSI and traditional networking, there are different types of cables of various lengths for use in a Fibre Channel configuration. Two types of cables are supported: copper and optical (fiber). Copper cables are used for short distances (up to 30m) and can be identified by their DB9 (9-pin) connector. Fiber cables come in two distinct types: Multi-Mode fiber (MMF) for short distances (up to 2 km), and Single-Mode Fiber (SMF) for longer distances (up to 10km). IBM will support the following distances shown in Table 6 for FCP.

Table 6. Cable and connector distances

Diameter(Microns)	Mode	Laser type	Distance
9	Single-mode	Longwave	<=10 km
50	Multi-mode	Shortwave	<=500 m
62.5	Multi-mode	Shortwave	<=175 m

In addition, connectors (see A.5.3, “Gigabit Interface Converters (GBIC)” on page 125) and cable adapters (see A.5.4, “Media Interface Adapters (MIA)” on page 125) have been developed that allow the interconnection of fiber optic based adapters with copper based devices.

A.5.2 Gigabit Link Model (GLM)

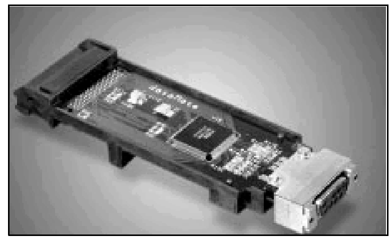

Gigabit Link Models are generic Fibre Channel transceiver units that integrate the key functions necessary for installation of a Fibre Channel media interface on most systems.

A.5.3 Gigabit Interface Converters (GBIC)

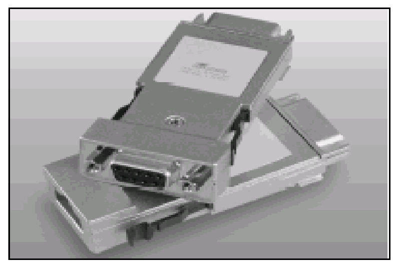

Gigabit Interface Converters are typically used with hubs and switches, and allow both copper and fiber optics to connect up to the same hub or switch. This works well in an environment where the components to be attached may be both fiber optics and copper, depending on cost and distance requirements.

A.5.4 Media Interface Adapters (MIA)

Media Interface Adapters can be used to facilitate conversion of an optical interface connection to copper and vice versa. Typically, MIAs are attached to host bus adaptors, but can also be used with switches and hubs. When a hub or switch only supports copper or optical connections, MIAs can be used to convert the signal to the appropriate media type (copper or optical).

A.5.5 Adapters

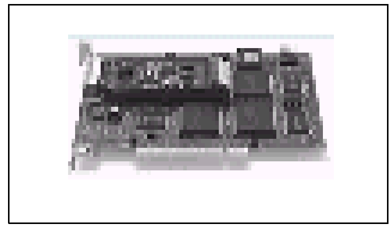

Adapters are devices that connect to a network, server, or storage device and control the electrical protocol for communications. Adapters are also referred to as Network Interface Cards (NIC), Enterprise Systems Connection (ESCON) adapters, Host Bus Adapters (HBA), and SCSI host bus adapters.

A.5.6 Extenders

Extenders are used to facilitate longer distances between nodes that exceed the theoretical maximum.

A.5.7 Multiplexors

Multiplexors provide for more effective utilization of high speed bandwidth resources by interleaving data from multiple sources onto a single Link. An example of this would be the FIbre CONnection Channel (FICON) Bridge, which allows up to eight separate ESCON paths to be multiplexed over a single FICON Link. The FICON bridge is not a separate box; as shown on the picture, it is a card in the ESCON Director. Multiplexors are also becoming increasingly efficient in terms of data compression, error correction, transmission speed and multi-drop capabilities.

A.5.8 Hubs

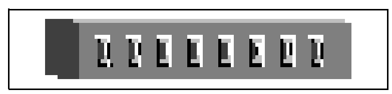

Fibre Channel hubs are used to connect up to 126 nodes into a logical loop. All connected nodes share the bandwidth of this one logical loop. Fibre Channel Arbitrated Loop (FC-AL) protocol is the most widely accepted, cost-effective alternative. Each port on a hub contains a Port Bypass Circuit (PBC) to automatically open and close the loop to support hot pluggability. Multiple hubs and links can be implemented to provide alternate path failover capability for high availability server environments. Intelligent hubs are currently being offered that provide features such as dynamic loop configuration and some of the benefits of switches.

A.5.9 Routers

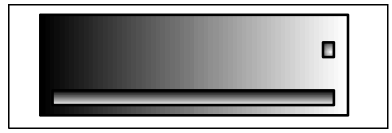

Storage routing is a new technology based on the old concept of routing as it is understood by the data communications industry. Storage routers differ from network routers in that the data being routed uses storage protocols like FCP (SCSI) instead of messaging protocols such as TCP/IP. The data path used to transfer storage data may be the same as that used for messaging traffic, but the content of the data itself contain imbedded storage protocol information. This is similar to the concept of tunneling protocols used in the broader market. For example, a storage router could encapsulate SCSI information in TCP/IP packets for transmission over an intervening ethernet network. The term *routing* implies data transfers over differing transmission media and addressing schemes. As a combination of communications and storage channel capabilities, Fibre Channel represents the first opportunity to apply communication techniques, such as routing, to storage traffic.

A.5.10 Bridges

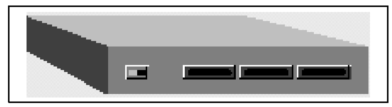

Bridges facilitate communication between LAN/SAN segments and/or networks with dissimilar protocols. An example of this would be a FICON bridge, which allows ESCON protocol to be tunneled over Fibre Channel protocol. FICON Bridges reduce the requirements of ESCON connections, ESCON channels, ESCON Director ports, and so on; they support large and small block multiplexing.

A.5.11 Gateways

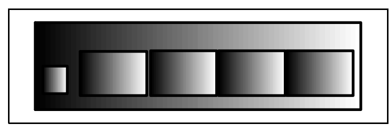

A gateway is a network station used to interconnect two or more dissimilar networks or devices, and may or may not perform protocol conversion. These boxes are typically used to provide access to WAN from a LAN. With gateways, SANs can be extended across a WAN.

Note: The IBM SAN Data Gateway is used to attach SCSI devices like tape libraries, disk subsystems at one end, and Fibre Channel connections at the other end, which makes it a router, not a gateway.

A.5.12 Switches

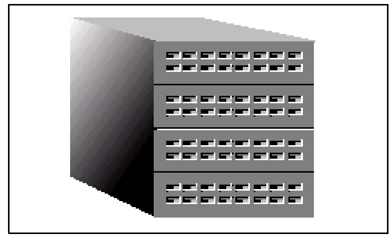

Switches are among the highest performing devices available for interconnecting large numbers of devices, increasing bandwidth, reducing congestion and providing aggregate throughput. Fibre Channel protocol was designed specifically by the computer industry to remove the barriers of performance with legacy channels and networks. When a Fibre Channel switch is implemented in a SAN, the network is referred to as a fabric, or switched fabric. Each device connected to a port on the switch can access any other device connected to any other port on the switch, enabling an on-demand connection to every connected device. Various FC switch offerings support both switched fabric and/or loop connections. As the number of devices increases, multiple switches can be cascaded for expanded access (fanout).

As switches allow any-to-any connections, the switch and management software can restrict which other ports a specific port can connect to. This is called port zoning.

A.5.13 Directors

When ESCON was announced, the switches were named *directors*, as they included high availability and other features; for example, dual power supply, not included in traditional switches.

A.5.13.1 ESCON

ESCON transfers data in half duplex mode at a transfer rate of 17 MB/sec. The directors provide dynamic or static connections between attached channels, control units, devices, and other ESCON directors. ESCON Directors provide single-link distances up to 3 km for multi-mode fiber and 20 km for single-mode fiber.

The ESCON director supports the FICON bridge card mentioned in A.5.10, "Bridges" on page 127. Using the FICON bridge, an S/390 system can, through a FICON channel, communicate with ESCON control units and devices attached to the director.

A.5.13.2 FICON

FICON transfers data in full duplex mode at a rate of 100 MB/sec. The FICON director is a regular Fibre Channel switch with some additional features. It provides connections between FICON channels and control units or devices with FICON interfaces. FICON Directors provide single-link distances up to 550 meters for multi-mode fiber (50 or 62.5 micron) and 10 km (20 km with RPQ) for single-mode fiber (9 micron).

Appendix B. Terminology and other basics

This Appendix explains a few terms used in the book and also gives explanations for SAN related topics.

B.1 StorWatch

StorWatch — IBM's Enterprise Storage Resource Management (ESRM) solution — is a growing software family whose goal is to enable storage administrators to efficiently manage storage resources from any location within an enterprise. Widely dispersed, disparate storage resources will ultimately be viewed and managed through a single, cohesive control point.

Agent

In the client-server model, this is the part of the system that performs information preparation and exchange on behalf of a client or server application. An example is the Agent SW installed on the SAN Data Gateway to talk to the StorWatch Data Gateway Specialist.

StorWatch Expert

These are StorWatch applications that employ a 3 tiered architecture that includes a management interface, a StorWatch manager, and agents that run on the storage resource(s) being managed. Expert products employ a StorWatch data base that can be used for saving key management data (for example, capacity or performance metrics). Expert products use the agents, as well as analysis of storage data saved in the data base, to perform higher value functions such as reporting of capacity, performance, and so on, over time (trends), configuration of multiple devices based on policies, monitoring of capacity and performance, automated responses to events or conditions, and storage related data mining.

StorWatch Reporter

StorWatch Reporter reduces disk management costs and increases the productivity of storage administrators.

StorWatch Reporter creates an inventory and collects disk capacity and utilization data for heterogeneous PC, UNIX, and OS/390 servers residing across your network. You can then generate customized reports that help you:

- Understand your present disk capacity and utilization
- Identify trends in your storage usage
- Anticipate and avoid outages due to out-of-storage conditions
- Plan for growth in your enterprise

Reporter also locates and launches Web based storage software applications for centralized administration of storage software and hardware.

StorWatch Specialist

A StorWatch interface for managing an individual Fibre Channel device or a limited number of like devices (that can be viewed as a single group). Storwatch specialists typically provide simple, point-in-time management functions such as configuration, reporting on asset and status information, simple device and event monitoring, and perhaps some service utilities.

B.2 SCSI

SCSI

Small Computer System Interface — A set of evolving ANSI standard electronic interfaces that allow personal computers to communicate with peripheral hardware such as disk drives, tape drives, CD_ROM drives, printers, and scanners, faster and more flexibly than previous interfaces. Table 7 below identifies the major characteristics of the different SCSI versions.

Table 7. Characteristics of different SCSI versions

SCSI Version	Signaling Rate (MHz)	Bus Width (bits)	Max. DTR (MB/s)	Max. Number Devices	Max. Cable Length
SCSI-1	5	8	5	7	6m
SCSI-2	5	8	5	7	6m
Wide SCSI-2	5	16	10	15	6m
Fast SCSI-2	10	8	10	7	6m
Fast Wide SCSI-2	10	16	20	15	6m
Ultra SCSI	20	8	20	7	1.5m
Ultra SCSI-2	20	16	40	7	12m
Ultra2 LVD SCSI	40	16	18	15	12m

SCSI-3

SCSI-3 consists of a set of primary commands and additional specialized command sets to meet the needs of specific device types. The SCSI-3 command sets are used not only for the SCSI-3 parallel interface but for additional parallel and serial protocols, including Fibre Channel, Serial Bus Protocol (used with IEEE 1394 Firewire physical protocol) and the Serial Storage Protocol (SSP).

SES

SCSI Enclosure Services — ANSI SCSI-3 proposal that defines a command set for soliciting basic device status (temperature, fan speed, power supply status, etc.) from a storage enclosures.

B.3 Fibre Channel terms

E-Port

Expansion Port — a port on a switch used to link multiple switches together into a Fibre Channel switch fabric.

F-Node

Fabric Node — a fabric attached node.

F-Port

Fabric Port — a port used to attach a Node Port (N_Port) to a switch fabric.

FL-Port

Fabric Loop Port — the access point of the fabric for physically connecting the user's Node Loop Port (NL_Port).

G-Port

Generic Port — a generic switch port that either a Fabric Port (F_Port) or an Expansion Port (E_Port) function is automatically determined during login.

L-Port

Loop Port — A node or fabric port capable of performing Arbitrated Loop functions and protocols. NL-Ports and FL_Ports are loop-capable ports.

N-Port

Node Port — A Fibre Channel-defined hardware entity at the end of a link which provides the mechanisms necessary to transport information units to or from another node.

NL-Port

Node Loop Port — a node port that supports Arbitrated Loop devices.

Point-to-Point Topology

An interconnection structure in which each point has physical links to only one neighbor resulting in a closed circuit. In point-to-point topology, the available bandwidth the is dedicated.

Zoning

In Fibre Channel environments, the grouping together of multiple ports to form a virtual private storage network. Ports that are members of a group or zone can communicate with each other but are isolated from ports in other zones.

Appendix C. Special notices

This publication is intended to help customers, Business Partners, and IBM technical professionals to understand and apply the Storage Area Network (SAN) concepts related to IBM tape library products. The information in this publication is not intended as the specification of any programming interfaces that are provided by described products. See the PUBLICATIONS section of the IBM Programming Announcement for each described product for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

This document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX
AT
ESCON
Magstar
OS/390
S/390
SP
System/390
400
AS/400
DFSMS
IBM
Netfinity
RS/6000
Seascape
StorWatch
Versatile Storage Server

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

SET and the SET logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Appendix D. Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 International Technical Support Organization publications

For information on ordering these ITSO publications see "How to get ITSO redbooks" on page 141.

- *Introduction to Storage Area Network, SAN*, SG24-5470
- *Tivoli Storage Manager Version 3.7: Technical Guide*, SG24-5477
- *Magstar MP 3575 Tape Library Dataserver: Multiplatform Implementation*, SG24-4983
- *IBM Magstar Tape Products Family: A Practical Guide*, SG24-4632

D.2 Redbooks on CD-ROMs

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at `http://www.redbooks.ibm.com/` for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	**Collection Kit Number**
System/390 Redbooks Collection	SK2T-2177
Networking and Systems Management Redbooks Collection	SK2T-6022
Transaction Processing and Data Management Redbooks Collection	SK2T-8038
Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
AS/400 Redbooks Collection	SK2T-2849
Netfinity Hardware and Software Redbooks Collection	SK2T-8046
RS/6000 Redbooks Collection (BkMgr)	SK2T-8040
RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

D.3 Other Publications

These publications are also relevant as further information sources:

- *IBM SAN Data Gateway Installation and User’s Guide*, SC26-7304
- *IBM SAN Fibre Channel Switch 2109 S16 User’s Guide*, SC26-7351
- *IBM SAN Fibre Channel Switch 2109 S08 User’s Guide*, SC26-7349

How to get ITSO redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** `http://www.redbooks.ibm.com/`

  Search for, view, download, or order hardcopy/CD-ROM redbooks from the redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this redbooks site.

  Redpieces are redbooks in progress; not all redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders by e-mail including information from the redbooks fax order form to:

	e-mail address
In United States	usib6fpl@ibmmail.com
Outside North America	Contact information is in the "How to Order" section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the "How to Order" section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the "How to Order" section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and redbooks by accessing the IBM Intranet Web site at `http://w3.itso.ibm.com/` and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at `http://w3.ibm.com/` for redbook, residency, and workshop announcements.

IBM Redbook Fax Order Form

Please send me the following:

Title	Order Number	Quantity

First name　Last name

Company

Address

City　Postal code　Country

Telephone number　Telefax number　VAT number

☐ Invoice to customer number

☐ Credit card number

Credit card expiration date　Card issued to　Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

T

U

V

W

Z

IBM Redbooks evaluation

Storage Area Networks: Tape Future in Fabrics
SG24-5474-00

Your feedback is very important to help us maintain the quality of IBM Redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at http://www.redbooks.ibm.com/
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?
_ Customer _ Business Partner _ Solution Developer _ IBM employee
_ None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction __________

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

SG24-5474-00
Printed in the U.S.A.

Storage Area Networks: Tape Future in Fabrics

SG24-5474-00