IBM

Planning and Implementing an IBM SAN

- Review planning considerations associated with Fibre Channel products
- Learn how to build an IBM SAN environment
- Set up and administer zoning

Jon Tate
George DeBiasi
Uwe Hofmann
Thomas Jahn

Redbooks

ibm.com/redbooks

SG24-6116-00

International Technical Support Organization

Planning and Implementing an IBM SAN

November 2000

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, “Special notices” on page 343.

First Edition (November 2000)

This edition applies to components, programs, architecture, and connections between multiple platforms and storage systems and a diverse range of software and hardware.

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 IBM Redbook is a follow-on from the redbook, *Designing an IBM Storage Area Network*, SG24-5758. In that book we introduced Fibre Channel basics, described the technical topology of a Storage Area Network (SAN), and detailed Fibre Channel products and IBM SAN initiatives. We also designed configurations that were able to maximize the benefits of Fibre Channel products that are currently supported by IBM and that are available in the marketplace today.

Where this IBM Redbook picks up the story is how to implement those products that are in the IBM product armory today. It is not possible to duplicate each and every SAN installation that is feasible or practical. What we want to achieve is a consolidated reference guide that details how the basic products can be swiftly and, in some cases, easily implemented. We will show the various features that each of these products contributes, and how the most common and important benefits of each are taken advantage of. We will show how they can be employed in some of the more commonly encountered environments and platforms.

With this in mind, we have two objectives within this redbook. The first is to show practical decisions to be considered when planning a SAN; the second objective is to show how the following products can be installed, configured, and tailored:

- IBM Enterprise Storage Server with native Fibre Channel
- IBM Fibre Channel Storage Hub
- IBM SAN Fibre Channel Managed Hub
- IBM SAN Fibre Channel Switch
- McDATA Enterprise Fibre Channel Director
- IBM Storage Area Network Data Gateway
- Vicom Fibre Channel SLIC Router

Once these products are successfully installed and all these configurations have been tested using a “hands-on” environment, we will show some of the benefits that we believe are fundamental to their application in a SAN.

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.

Jon Tate is a project leader for SAN Storage Solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level-2 support for IBM storage products. Jon has 14 years of experience in storage software and management, services, and support.

George DeBiasi is an IT Specialist in the Americas Storage Techline in Toronto. Before joining the Techline in 1998, he provided pre-sales and technical support on videoconferencing systems for IBM and its customers. George has 14 years of experience as a Technical Specialist providing customer service, support and project management. He holds a diploma in Electronic Engineering Technology from Radio College of Canada.

Uwe Hofmann is a Senior Consultant and Sales Support Manager in IBM's Europe, Middle East and Africa (EMEA) SAN and Storage Solutions Center team. He joined IBM more than 28 years ago in the days of System/360, and is now enjoying the frenetic world of SAN. He has 13 years experience in IBM's storage business. He held a number of sales and marketing roles in different branch offices in Germany. Since 1988 he has been based in the EMEA SSD Customer Executive Briefing Center in Mainz, Germany. Uwe holds a diploma in National Economics. He is a regular speaker on SAN and storage related topics for IBM customer groups and external conferences in Europe.

Thomas Jahn is an IT Specialist within the IBM Storage Subsystems Division in Germany. He has three years of experience providing technical support in IBM. Thomas has provided technical support for networking and server consolidation on OS/390 UNIX for IBM and its customers. He is currently engaged in providing support for open systems storage across multiple platforms and a wide customer base. He holds a Dipl. Ing. degree in Computer Science from the Staatliche Studienakademie Sachsen.

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

Lisa Haut-Mikkelsen
IBM Storage Subsystems Division

Robert Azevedo
IBM Storage Subsystems Division

Jim Baldyga
Brocade

Mark Bruni
IBM Storage Subsystems Division

Joe Emery
Vicom Systems, Inc.

Glenda Fuller
IBM Storage Subsystems Division

Malkit Hayun
IBM Storage Subsystems Division

Ronda Hruby
IBM Storage Subsystems Division

Scott Jensen
Brocade

Chris Morton
IBM Storage Subsystems Division

Dietmar Kurpanek
Vicom Systems, Inc.

Paul Radu
Pathlight Technology, Inc.

Tammy Sokol
IBM Storage Subsystems Division

Karen Ward
IBM Storage Subsystems Division

John Young
IBM Storage Subsystems Division

Norm Weinberg
Vicom Systems, Inc.

Sven Eichelbaum
IBM Storage Subsystems Division

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 review” on page 371 to the fax number shown on the form.
- Use the online evaluation form found at **ibm.com**/redbooks
- Send your comments in an Internet note to redbook@us.ibm.com

Chapter 1. Introduction to Storage Area Networks

Everyone working in the Information Technology industry is familiar with the continuous developments in technology, which constantly deliver improvements in performance, capacity, size, functionality and so on. A few of these developments have far reaching implications because they enable applications or functions which allow us fundamentally to rethink the way we do things and go about our everyday business. The advent of Storage Area Networks (SANs) is one such development. SANs can lead to a proverbial "paradigm shift" in the way we organize and use the IT infrastructure of an enterprise.

In the chapter that follows, we show the market forces that have driven the need for a new storage infrastructure, coupled with the benefits that a SAN brings to the enterprise.

1.1 The need for a new storage infrastructure

The 1990's witnessed a major shift away from the traditional mainframe, host-centric model of computing to the client/server model. Today, many organizations have hundreds, even thousands, of distributed servers and client systems installed throughout the enterprise. Many of these systems are powerful computers, with more processing capability than many mainframe computers had only a few years ago.

Storage, for the most part, is directly connected by a dedicated channel to the server it supports. Frequently the servers are interconnected using local area networks (LAN) and wide area networks (WAN), to communicate and exchange data. This is illustrated in Figure 1. The amount of disk storage capacity attached to such systems has grown exponentially in recent years. It is commonplace for a desktop Personal Computer (PC) today to have 5 gigabytes (GB) or 10 gigabytes, and single disk drives with up to 75 GB are available. There has been a move to disk arrays, comprising a number of disk drives. The arrays may be "just a bunch of disks" (JBOD), or various implementations of redundant arrays of independent disks (RAID). The capacity of such arrays may be measured in tens or hundreds of GBs, but I/O bandwidth has not kept pace with the rapid growth in processor speeds and disk capacities.

Distributed clients and servers are frequently chosen to meet specific application needs. They may, therefore, run different operating systems (such as Windows NT, UNIX of differing flavors, Novell Netware, VMS and so on), and different database software (for example, DB2, Oracle, Informix, SQL

Server). Consequently, they have different file systems and different data formats.

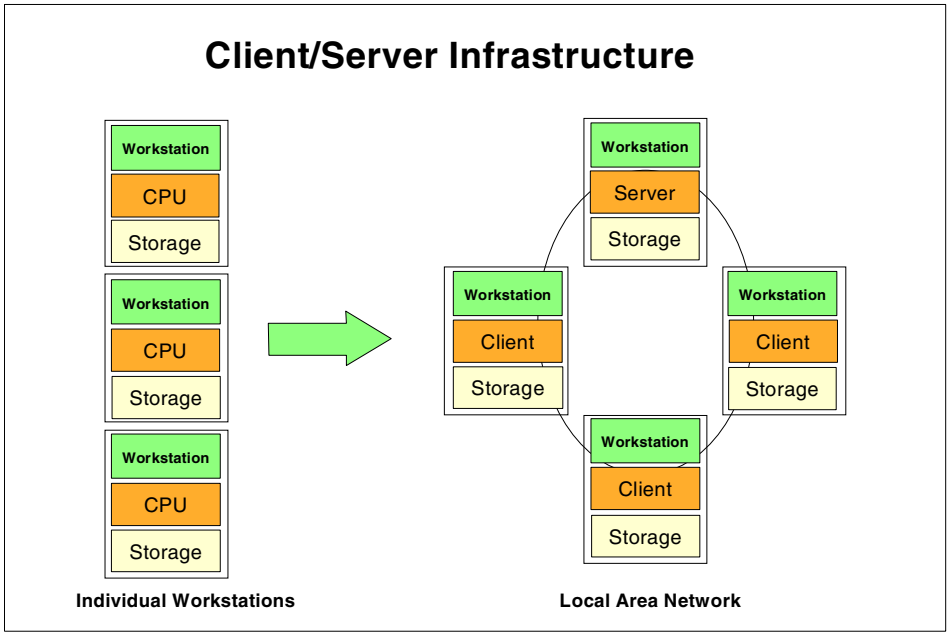

Figure 1. Typical distributed systems or client server infrastructure

Managing this multi-platform, multi-vendor, networked environment has become increasingly complex and costly. Multiple vendor's software tools, and appropriately-skilled human resources must be maintained to handle data and storage resource management on the many differing systems in the enterprise. Surveys published by industry analysts consistently show that management costs associated with distributed storage are much greater, up to 10 times more, than the cost of managing consolidated or centralized storage. This includes costs of backup, recovery, space management, performance management and disaster recovery planning.

Disk storage is often purchased from the processor vendor as an integral feature, and it is difficult to establish if the price you pay per gigabyte is competitive, compared to the market price of disk storage. Disks and tape drives, directly attached to one client or server, cannot be used by other systems, leading to inefficient use of hardware resources. Organizations often find that they have to purchase more storage capacity, even though free capacity is available, but is attached to other platforms. This is illustrated in Figure 2.

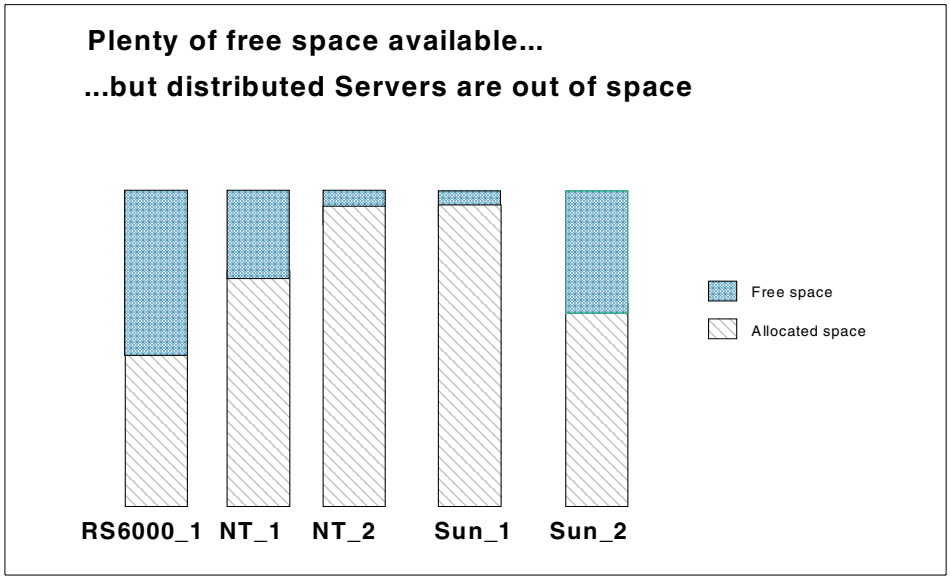

Figure 2. Inefficient use of available disk capacity attached to individual servers

Additionally, it is difficult to scale capacity and performance to meet rapidly changing requirements, such as the explosive growth in e-business applications.

Data stored on one system cannot readily be made available to other users, except by creating duplicate copies, and moving the copy to storage that is attached to another server. Movement of large files of data may result in significant degradation of performance of the LAN/WAN, causing conflicts with mission critical applications. Multiple copies of the same data may lead to inconsistencies between one copy and another. Data spread on multiple small systems is difficult to coordinate and share for enterprise-wide applications, such as e-business, Enterprise Resource Planning (ERP), Data Warehouse, and Business Intelligence (BI).

Backup and recovery operations across a LAN may also cause serious disruption to normal application traffic. Even using fast Gigabit Ethernet transport, sustained throughput from a single server to tape is about 25 GB per hour. It would take approximately 12 hours to fully backup a relatively moderate departmental database of 300 GBs. This may exceed the available window of time in which this must be completed, and it may not be a practical solution if business operations span multiple time zones. It is increasingly evident to IT managers that these characteristics of client/server computing are too costly, and too inefficient. The islands of information resulting from the

distributed model of computing do not match the needs of the e-business enterprise. We show this in Figure 3.

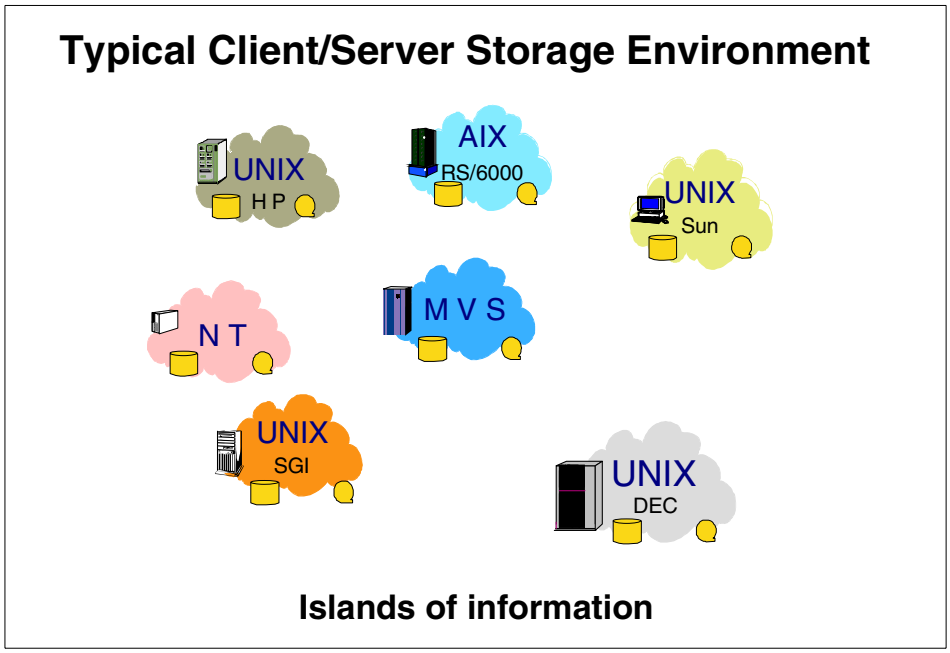

Figure 3. Distributed computing model tends to create islands of information

New ways must be found to control costs, to improve efficiency, and to properly align the storage infrastructure to meet the requirements of the business. One of the first steps to improved control of computing resources throughout the enterprise is improved connectivity.

In the topics that follow, we look at the advantages and disadvantages of the standard storage infrastructure of today.

1.2 The Small Computer Systems Interface legacy

The Small Computer Systems Interface (SCSI) is the conventional, server centric method of connecting peripheral devices (disks, tapes and printers) in the open client/server environment. As its name indicates, it was designed for the PC and small computer environment. It is a bus architecture, with dedicated, parallel cabling between the host and storage devices, such as disk arrays. This is similar in implementation to the Original Equipment Manufacturer's Information (OEMI) bus and tag interface commonly used by mainframe computers until the early 1990's. SCSI shares a practical aspect

with bus and tag, in that cables and connectors are bulky, relatively expensive, and are prone to failure.

The amount of data available to the server is determined by the number of devices which can attach to the bus, and by the number of buses attached to the server. Up to 15 devices can be attached to a server on a single SCSI bus. In practice, because of performance limitations due to arbitration, it is common for no more than four or five devices to be attached in this way, therefore limiting capacity scalability.

Access to data is lost in the event of a failure of any of the SCSI connections to the disks. This also applies in the event of reconfiguration or servicing of a disk device attached to the SCSI bus, because all the devices in the string must be taken offline. In today's environment, when many applications need to be available continuously, this downtime is unacceptable.

The data rate of the SCSI bus is determined by the number of bits transferred, and the bus cycle time (measured in megahertz (MHz)). Decreasing the cycle time increases the transfer rate, but, due to limitations inherent in the bus architecture, it may also reduce the distance over which the data can be successfully transferred. The physical transport was originally a parallel cable comprising eight data lines, to transmit eight bits in parallel, plus control lines. Later implementations widened the parallel data transfers to 16 bit paths (SCSI Wide), to achieve higher bandwidths.

Propagation delays in sending data in parallel along multiple lines lead to a well known phenomenon known as skew, meaning that all bits may not arrive at the target device at the same time. This is shown in Figure 4.

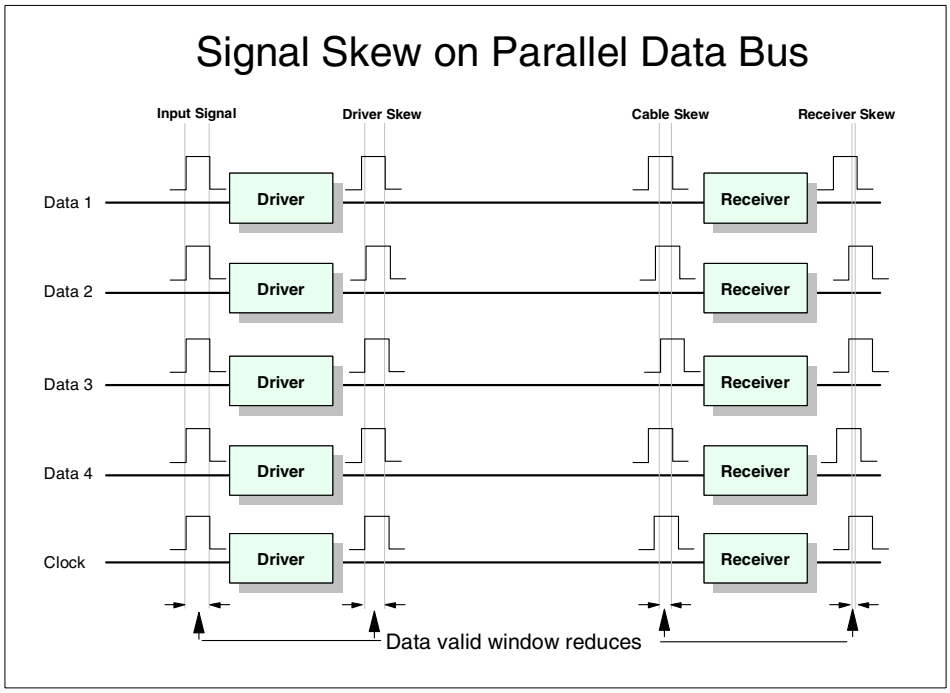

Figure 4. SCSI propagation delay results in skew

Arrival occurs during a small window of time, depending on the transmission speed, and the physical length of the SCSI bus. The need to minimize the skew limits the distance that devices can be positioned away from the initiating server to between 2 meters to 25 meters, depending on the cycle time. Faster speed means shorter distance. The distances refer to the maximum length of the SCSI bus, including all attached devices. The SCSI distance limitations are shown in Figure 5. These distance limitations may severely restrict the total GB capacity of the disk storage which can be attached to an individual server.

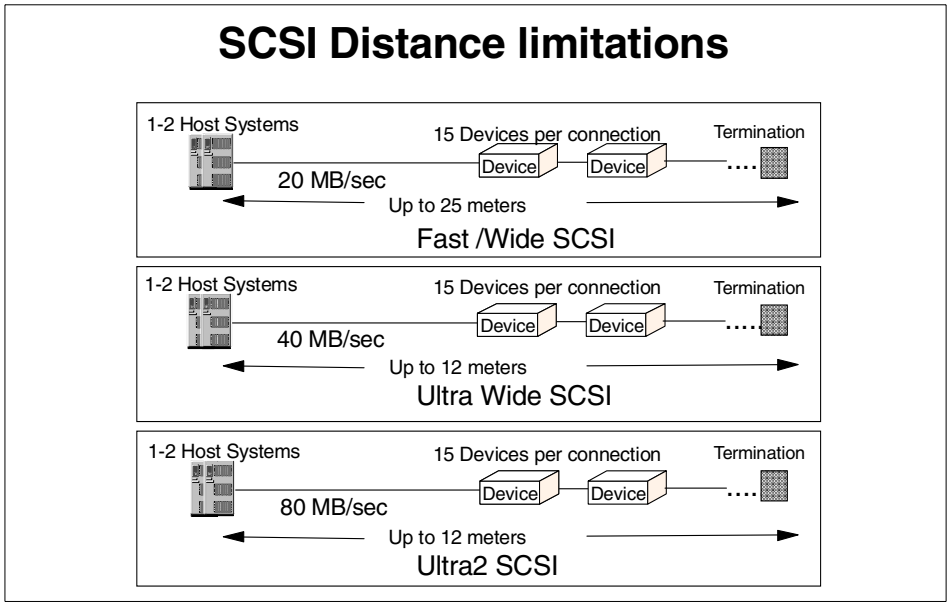

Figure 5. SCSI bus distance limitations

Many applications require the system to access several devices, or for several systems to share a single device. SCSI can enable this by attaching multiple servers or devices to the same bus. This is known as a multi-drop configuration. A multi-drop configuration is shown in Figure 6.

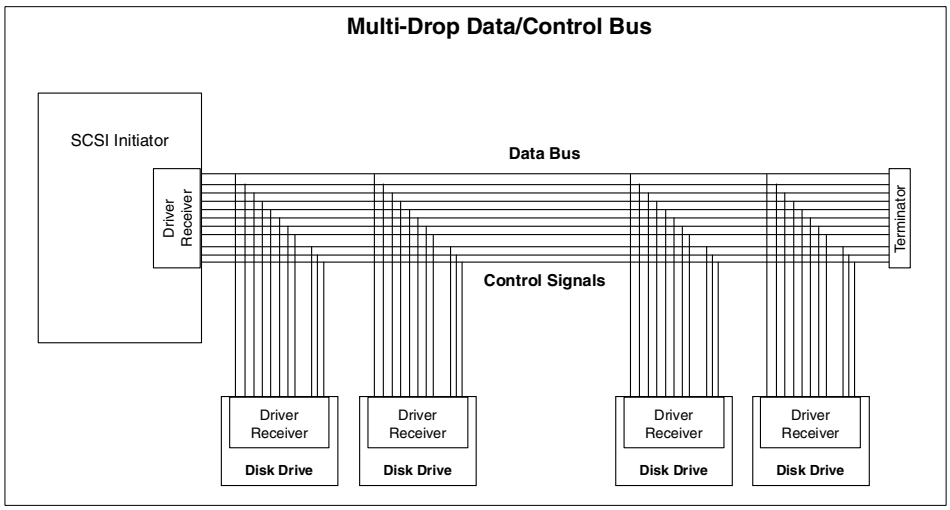

Figure 6. Multi-drop bus structure

To avoid signal interference, and therefore possible data corruption, all unused ports on a parallel SCSI bus must be properly terminated. Incorrect termination can result in transaction errors or failures.

Normally, only a single server can access data on a specific disk by means of a SCSI bus. In a shared bus environment, it is clear that all devices cannot transfer data at the same time. SCSI uses an arbitration protocol to determine which device can gain access to the bus. Arbitration occurs before and after every data transfer on the bus. While arbitration takes place, no data movement can occur. This represents an additional overhead which reduces bandwidth utilization, substantially reducing the effective data rate achievable on the bus. Actual rates are typically less than 50% of the rated speed of the SCSI bus.

In addition to being a physical transport, SCSI is also a protocol, which specifies commands and controls for sending blocks of data between the host and the attached devices. SCSI commands are issued by the host operating system, in response to user requests for data. Some operating systems, for example, Windows NT, treat all attached peripherals as SCSI devices and issue SCSI commands to deal with all read and write operations.

It is clear that the physical parallel SCSI bus architecture has a number of significant speed, distance, and availability limitations, which make it increasingly less suitable for many applications in today's networked IT infrastructure. However, since the SCSI protocol is deeply embedded in the way that commonly encountered operating systems handle user requests for data, it would be a major inhibitor to progress if we were obliged to move to new protocols.

1.3 Storage network solutions

Today's enterprise IT planners need to link many users of multi-vendor, heterogeneous systems to multi-vendor shared storage resources, and they need to allow those users to access common data, wherever it is located in the enterprise. These requirements imply a network solution, and two types of network storage solutions are now available:

- Network attached storage (NAS)
- Storage Area Network (SAN)

1.3.1 What network attached storage is

NAS solutions utilize the LAN in front of the server, and transmit data over the LAN using messaging protocols, such as TCP/IP and Net BIOS. We illustrate this in Figure 7.

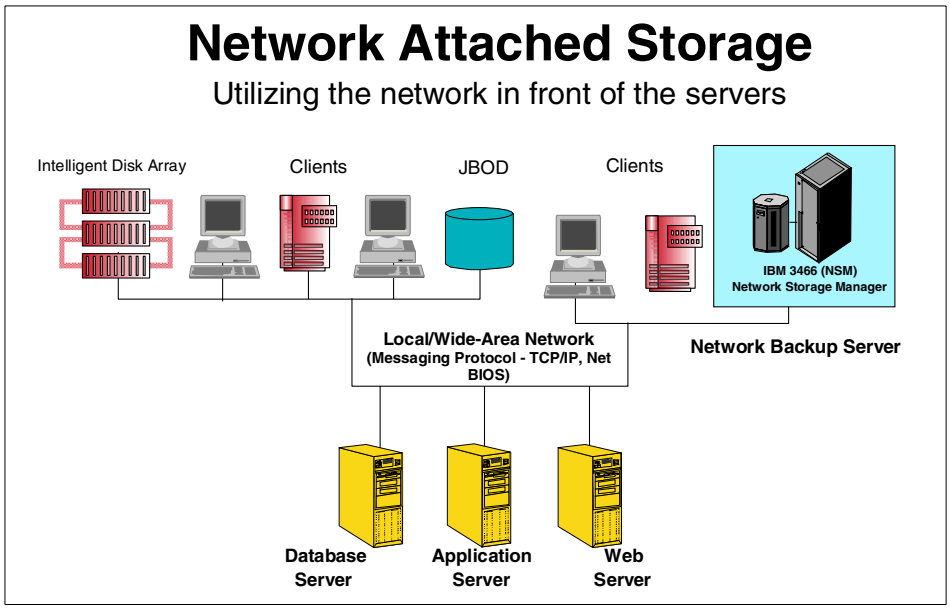

Figure 7. Network attached storage - utilizing the network in front of the servers

By making storage devices LAN addressable, the storage is freed from its direct attachment to a specific server. In principle, any user running any operating system can address the storage device by means of a common access protocol, for example, Network File System (NFS). In addition, a task, such as back-up to tape, can be performed across the LAN, enabling sharing of expensive hardware resources between multiple servers. Most storage devices cannot just attach to a LAN. NAS solutions are specialized file servers which are designed for this type of attachment.

NAS, therefore, offers a number of benefits, which address some of the limitations of parallel SCSI. However, by moving storage transactions, such as disk accesses, and tasks, such as backup and recovery of files, to the LAN, conflicts can occur with end user traffic on the network. LANs are tuned to favor short burst transmissions for rapid response to messaging requests, rather than large continuous data transmissions. Significant overhead can be imposed to move large blocks of data over the LAN, due to the small packet

size used by messaging protocols. For instance, the maximum packet size for Ethernet is about 1500 bytes. A 10 MB file has to be segmented into more than 7000 individual packets, (each sent separately by the LAN access method), if it is to be read from a NAS device. Therefore, a NAS solution is best suited to handle cross platform direct access applications, not to deal with applications requiring high bandwidth.

NAS solutions are relatively low cost, and straightforward to implement as they fit in to the existing LAN environment, which is a mature technology. However, the LAN must have plenty of spare capacity to justify NAS implementations. A number of vendors, including IBM, offer a variety of NAS solutions. These fall into two categories:

- File servers
- Backup/archive servers

However, it is not the purpose of this book to discuss these. NAS can be used separately or together with a SAN, as the technologies are complementary. In general terms, NAS offers lower cost solutions, but with more limited benefits, lower performance, and less scalability, than Fibre Channel SANs.

1.3.2 What a Storage Area Network is

A SAN is a specialized, high speed network attaching servers and storage devices. It is sometimes called “the network behind the servers”. A SAN allows “any to any” connection across the network, using interconnect elements such as routers, gateways, hubs and switches. It eliminates the traditional dedicated connection between a server and storage, and the concept that the server effectively “owns and manages” the storage devices. It also eliminates any restriction to the amount of data that a server can access, currently limited by the number of storage devices, which can be attached to the individual server. Instead, a SAN introduces the flexibility of networking to enable one server or many heterogeneous servers to share a common storage “utility”, which may comprise many storage devices, including disk, tape, and optical storage. And, the storage utility may be located far from the servers which use it. We show what the network behind the servers may look like, in Figure 8.

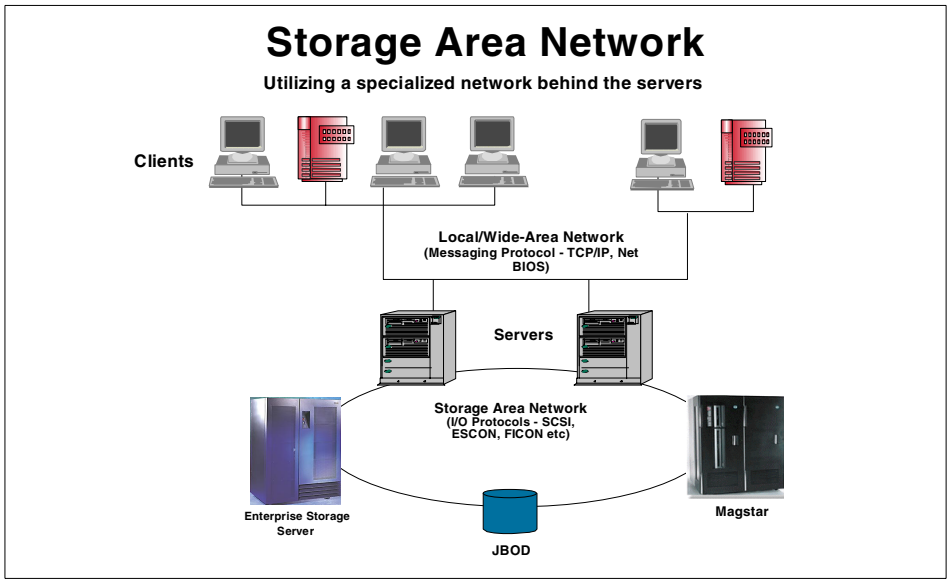

Figure 8. Storage Area Network - the network behind the servers

A SAN differs from traditional networks, because it is constructed from storage interfaces. SAN solutions utilize a dedicated network behind the servers, based primarily (though, not necessarily) on Fibre Channel architecture. Fibre Channel provides a highly scalable bandwidth over long distances, and with the ability to provide full redundancy, including switched, parallel data paths to deliver high availability and high performance.

Therefore, a SAN can bypass traditional network bottlenecks. It supports direct, high speed transfers between servers and storage devices in the following ways:

- **Server to storage**: This is the traditional method of interaction with storage devices. The SAN advantage is that the same storage device may be accessed serially or concurrently by multiple servers.
- **Server to server**: This is high speed, high volume communications between servers.
- **Storage to storage**: For example, a disk array can back up its data direct to tape across the SAN, without processor intervention. Or, a device can be mirrored remotely across the SAN.

A SAN changes the server centric model of the typical open systems IT infrastructure, replacing it with a data centric infrastructure.

1.3.3 What about ESCON and FICON?

Sceptics might already be saying that the concept of SAN is not new. Indeed, for System 390 (S/390) users, the implementation of shared storage on a dedicated network has been common since the introduction of Enterprise System Connection (ESCON) in 1991.

However, for UNIX, Windows NT, and other open systems users, the need for such capability is now extremely high. As we have shown, the traditional SCSI parallel bus architecture, most commonly used in these environments, is no longer capable of handling their growing data intensive application requirements. These users are faced with many of the same problems which challenged mainframe users in the late 1980s and early 1990s, and which largely were solved by ESCON.

But the ESCON architecture does not answer the open systems needs of today, due to a number of critical limitations. ESCON is primarily a S/390 solution, which does not support open systems protocols for data movement, and ESCON is limited in performance (nominally 17 MB/second), relative to technologies available today. An enhancement to ESCON is provided by Fibre Connection (FICON). Figure 9 shows how FICON enhances ESCON.

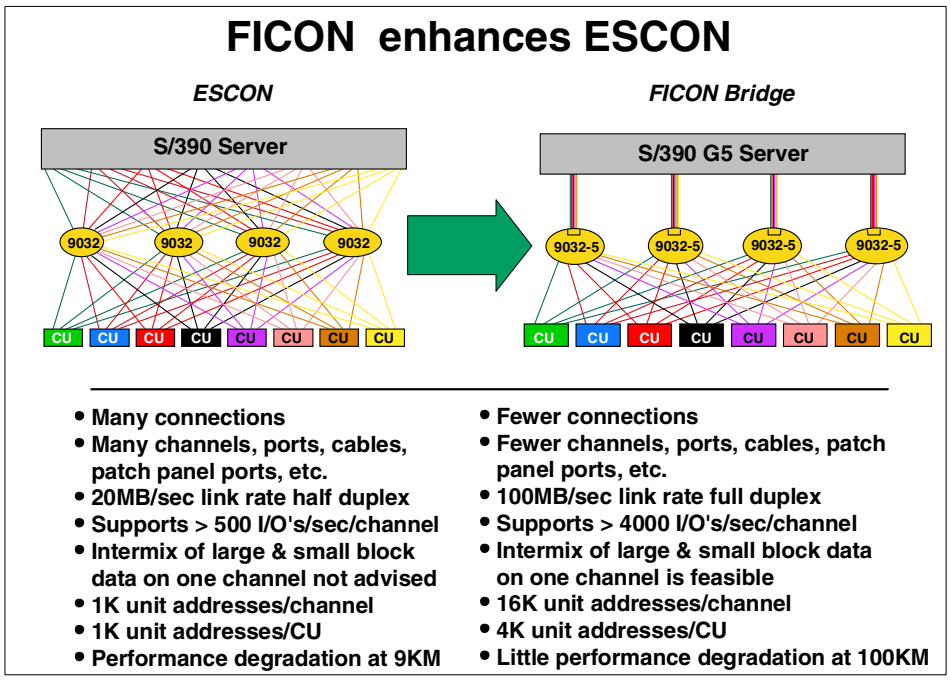

Figure 9. FICON enhances ESCON

The S/390 FICON architecture retains ESCON topology and switch management characteristics. FICON channels can deliver data rates up to 100 MB/second full-duplex, and they extend channel distances up to 100 kilometers. More storage controllers and devices can be supported per FICON link, relieving channel constraints in configuring S/390 processors.

The FICON architecture is fully compatible with existing S/390 channel command words (CCWs) and programs. But, most importantly, FICON uses Fibre Channel for transportation of data, and, therefore in principle, is capable of participating with other platforms (UNIX, Windows NT, Novell Netware, and so on) in a Fibre Channel enterprise SAN. However, this capability is not yet supported, due to a number of network management requirements imposed by the S/390 architecture.

IBM expects a transition period, during which S/390 FICON SANs will develop separately from Fibre Channel Protocol (FCP) open systems SANs, which use the SCSI protocol. In the longer term, FCP and FICON SANs will merge into a true Enterprise SAN. IBM has published a number of redbooks on the subject of FICON and an example of this is *Introduction to IBM S/390 FICON*, SG24-5176. Additional redbooks that describe FICON can be found at the IBM Redbooks site by using the search argument *FICON*.

`http://www.ibm.com/redbooks`

For this reason, this book will focus exclusively on FCP open systems elements of IBM's Enterprise SAN which are available today.

1.4 What Fibre Channel is

Fibre Channel is an open, technical standard for networking which incorporates the "channel transport" characteristics of an I/O bus, with the flexible connectivity and distance characteristics of a traditional network. Notice the European spelling of Fibre, which is intended to distinguish it from fiber-optics and fiber-optic cabling, which are physical hardware and media used to transmit data at high speed over long distances using light emitting diode (LED) and laser technology.

Because of its channel-like qualities, hosts and applications see storage devices attached to the SAN as if they are locally attached storage. Because of its network characteristics it can support multiple protocols and a broad range of devices, and it can be managed as a network. Fibre Channel can use either optical fiber (for distance) or copper cable links (for short distance at low cost).

Fibre Channel is a multi-layered network, based on a series of American National Standards Institute (ANSI) standards, which define characteristics and functions for moving data across the network. These include definitions of physical interfaces, such as cabling, distances and signaling; data encoding and link controls; data delivery in terms of frames, flow control and classes of service; common services; and protocol interfaces.

Like other networks, information is sent in structured packets or frames, and data is serialized before transmission. But, unlike other networks, the Fibre Channel architecture includes a significant amount of hardware processing to deliver high performance. The speed currently achieved is 100 MB per second, (with the potential for 200 MB and 400 MB and higher data rates in the future). In all Fibre Channel topologies, a single transmitter sends information to a single receiver. In most multi-user implementations this requires that routing information (source and target) must be provided. Transmission is defined in the Fibre Channel standards across three transport topologies:

- **Point to point**: A bidirectional, dedicated interconnection between two nodes, with full-duplex bandwidth (100 MB/second in each direction concurrently).
- **Arbitrated loop**: A unidirectional ring topology, similar to a token ring, supporting up to 126 interconnected nodes. Each node passes data to the next node in the loop, until the data reaches the target node. All nodes share the 100 MB/second Fibre Channel bandwidth. Devices must arbitrate for access to the loop. Therefore, with 100 active devices on a loop, the effective data rate for each is 1 MB/second, which is further reduced by the overhead of arbitration. A loop may also be connected to a Fibre Channel switch port, therefore enabling attachment of the loop to a wider switched fabric environment. In this case, the loop may support up to 126 devices.

  Many fewer devices are normally attached in practice, because of arbitration overheads and shared bandwidth constraints. Due to fault isolation issues inherent with arbitrated loops, most Fibre Channel arbitrated loop (FC-AL) SANs have been implemented with a maximum of two servers, plus a number of peripheral storage devices. So FC-AL is suitable for small SAN configurations, or SANlets.
- **Switched fabric**: The term *fabric* describes an intelligent switching infrastructure which delivers data from any source to any destination. The interconnection of up to 2^{24} nodes is allowed, with each node able to utilize the full 100 MB/second duplex Fibre Channel bandwidth. Each logical connection receives dedicated bandwidth, so the overall bandwidth is multiplied by the number of connections (delivering a maximum of 200

MB/second x *n* nodes). The fabric itself is responsible for controlling the routing of information. It may be simply a single switch, or it may comprise multiple interconnected switches which function as a single logical entity. Complex fabrics must be managed by software which can exploit SAN management functions which are built into the fabric. Switched fabric is the basis for enterprise-wide SANs.

A mix of these three topologies can be implemented to meet specific needs. Fibre Channel Arbitrated Loop (FC-AL) and switched fabric (FC-SW) are the two most commonly used topologies, satisfying differing requirements for scalability, distance, cost and performance. A fourth topology has been developed, known as slotted loop (FC-SL); But, this appears to have limited application, specifically in aerospace, so it is not discussed in this book.

Fibre Channel uses a serial data transport scheme, similar to other computer networks, streaming packets, (frames) of bits one behind the other in a single data line. To achieve the high data rate of 100 MB/second the transmission clock frequency is currently one gigabit, or one bit per 0.94 nanoseconds.

Serial transfer, of course, does not suffer from the problem of skew, so speed and distance is not restricted as with parallel data transfers as we show in Figure 10.

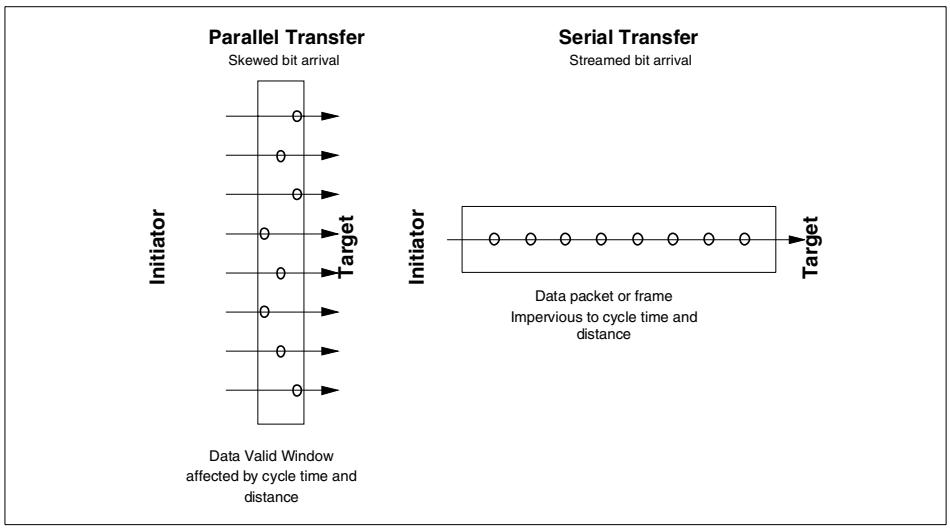

Figure 10. Parallel data transfers versus serial data transfers

Serial transfer enables simpler cabling and connectors, and also routing of information through switched networks. Today, Fibre Channel can operate over distances of up to 10 km, link distances up to 90 km by implementing cascading, and longer with the introduction of repeaters. Just as LANs can be interlinked in WANs by using high speed gateways, so can campus SANs be interlinked to build enterprise-wide SANs.

Whatever the topology, information is sent between two nodes, which are the source (transmitter or initiator) and destination (receiver or target). A node is a device, such as a server (personal computer, workstation, or mainframe), or peripheral device, such as disk or tape drive, or video camera. Frames of information are passed between nodes, and the structure of the frame is defined by a protocol. Logically, a source and target node must utilize the same protocol, but each node may support several different protocols or data types.

Therefore, Fibre Channel architecture is extremely flexible in its potential application. Fibre Channel transport layers are protocol independent, enabling the transmission of multiple protocols. It is possible, therefore, to transport storage I/O protocols and commands, such as SCSI-3 Fibre Channel Protocol, (or FCP, the most common implementation today), ESCON, FICON, SSA, and HIPPI. Network packets may also be sent using messaging protocols, for instance, TCP/IP or Net BIOS, over the same physical interface using the same adapters, cables, switches, and other infrastructure hardware. Theoretically, then multiple protocols can move concurrently over the same fabric. This capability is not in common use today, and, in any case, currently excludes concurrent FICON transport. Most Fibre Channel SAN installations today only use a single protocol.

Using a credit based flow control methodology, Fibre Channel is able to deliver data as fast as the destination device buffer is able to receive it. And, low transmission overheads enable high sustained utilization rates without loss of data.

Therefore, Fibre Channel combines the best characteristics of traditional I/O channels with those of computer networks:

- High performance for large data transfers by using simple transport protocols and extensive hardware assists
- Serial data transmission
- A physical interface with a low error rate definition
- Reliable transmission of data with the ability to guarantee or confirm error-free delivery of the data

- Packaging data in packets (*frames* in Fibre Channel terminology)
- Flexibility in terms of the types of information which can be transported in frames (such as data, video and audio)
- Use of existing device-oriented command sets, such as SCSI and FCP
- A vast expansion in the number of devices which can be addressed when compared to I/O interfaces — a theoretical maximum of more than 16 million ports

It is this high degree of flexibility, availability, and scalability; the combination of multiple protocols at high speeds over long distances; and the broad acceptance of the Fibre Channel standards by vendors throughout the IT industry, which makes the Fibre Channel architecture ideal for the development of enterprise SANs.

For more details of the Fibre Channel architecture, refer to *Designing an IBM Storage Area Network*, SG24-5758.

1.5 The business benefits of a Fibre Channel SAN

Today's business environment creates many challenges for the enterprise IT planner. SANs can provide solutions to many of their operational problems.

1.5.1 Storage consolidation and sharing of resources

By enabling storage capacity to be connected to servers at a greater distance, and by disconnecting storage resource management from individual hosts, a SAN enables disk storage capacity to be consolidated. The results can be lower overall costs through better utilization of the storage, lower management costs, increased flexibility, and increased control.

This can be achieved physically or logically.

1.5.1.1 Physical consolidation

Data from disparate storage subsystems can be combined on large, enterprise-class, shared disk arrays, which may be located at some distance from the servers. The capacity of these disk arrays can be shared by multiple servers, and users may also benefit from the advanced functions typically offered with such subsystems. This may include RAID capabilities, remote mirroring, and instantaneous data replication functions, which might not be available with smaller, integrated disks.The array capacity may be partitioned, so that each server has an appropriate portion of the available GBs. This is shown in Figure 11.

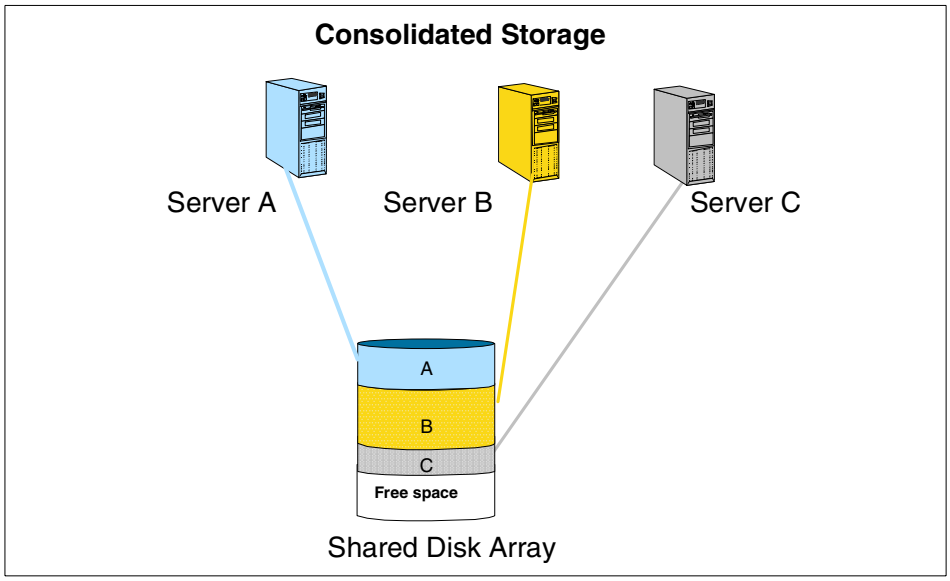

Figure 11. Consolidated storage - efficiently shared capacity

Available capacity can be dynamically allocated to any server requiring additional space. Capacity not required by a server application can be reallocated to other servers. This avoids the inefficiency associated with free disk capacity attached to one server not being usable by other servers. Extra capacity may be added, in a non-disruptive manner.

1.5.1.2 Logical consolidation

It is possible to achieve shared resource benefits from the SAN, but without moving existing equipment. A SAN relationship can be established between a client and a group of storage devices that are not physically co-located (excluding devices which are internally attached to servers). A logical view of the combined disk resources may allow available capacity to be allocated and reallocated between different applications running on distributed servers, to achieve better utilization. Consolidation is covered in greater depth in *IBM Storage Solutions for Server Consolidation*, SG24-5355.

Figure 12 shows a logical consolidation of independent arrays.

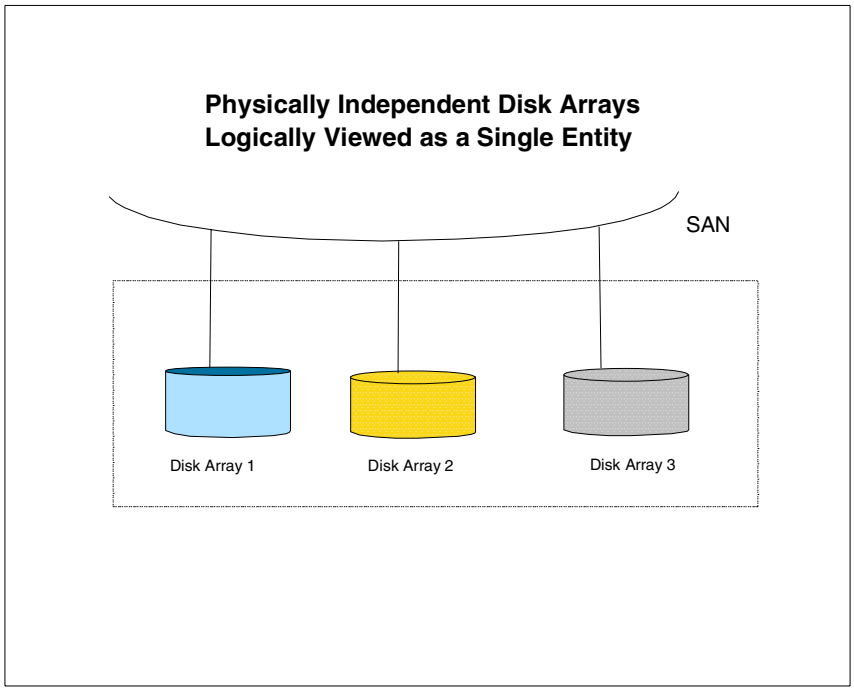

Figure 12. Logical consolidation of dispersed disk subsystems

1.5.2 Data sharing

The term *data sharing* is used somewhat loosely by users and some vendors. It is sometimes interpreted to mean the replication of files or databases to enable two or more users, or applications, to concurrently use separate copies of the data. The applications concerned may operate on different host platforms. A SAN may ease the creation of such duplicated copies of data using facilities such as remote mirroring.

Data sharing may also be used to describe multiple users accessing a single copy of a file. This can be called *true data sharing*. In a homogeneous server environment, with appropriate application software controls, multiple servers may access a single copy of data stored on a consolidated storage subsystem.

If attached servers are heterogeneous platforms (for example, a mix of UNIX and Windows NT), sharing of data between such unlike operating system environments is complex. This is due to differences in file systems, data formats, and encoding structures. IBM, however, uniquely offers a true data

sharing capability, with concurrent update, for selected heterogeneous server environments, using the Tivoli SANergy File Sharing solution. The details can be found at:

http://www.sanergy.com

The SAN advantage in enabling enhanced data sharing may reduce the need to hold multiple copies of the same file or database. This reduces duplication of hardware costs to store such copies. It also enhances the ability to implement cross enterprise applications, such as e-business, which may be inhibited when multiple data copies are stored.

1.5.3 Non-disruptive scalability for growth

There is an explosion in the quantity of data stored by the majority of organizations. This is fueled by the implementation of applications, such as e-business, e-mail, Business Intelligence, Data Warehouse, and Enterprise Resource Planning. Industry analysts, such as external storage consulting groups, estimate that electronically stored data is doubling every year. In the case of e-business applications, opening the business to the Internet, there have been reports of data growing by more than 10 times annually. This is a nightmare for planners, as it is increasingly difficult to predict storage requirements.

A finite amount of disk storage can be connected physically to an individual server due to adapter, cabling and distance limitations. With a SAN, new capacity can be added as required, without disrupting ongoing operations. SANs enable disk storage to be scaled independently of servers.

1.5.4 Improved backup and recovery

With data doubling every year, what effect does this have on the backup window? Back-up to tape, and recovery, are operations which are problematic in the parallel SCSI or LAN based environments. For disk subsystems attached to specific servers, two options exist for tape backup. Either it must be done to a server attached tape subsystem, or by moving data across the LAN.

1.5.4.1 Tape pooling

Providing tape drives to each server is costly, and it also involves the added administrative overhead of scheduling the tasks and managing the tape media. SANs allow for greater connectivity of tape drives and tape libraries, especially at greater distances. Tape pooling is the ability for more than one server to logically share tape drives within an automated library. This can be achieved by software management, by using tools, such as Tivoli Storage

Manager, or with tape libraries with outboard management, such as IBM's 3494.

1.5.4.2 LAN-free and server-free data movement

Backup using the LAN moves the administration to centralized tape drives or automated tape libraries. However, at the same time, the LAN experiences very high traffic volume during the backup or recovery operations, and this can be extremely disruptive to normal application access to the network. Although backups can be scheduled during non-peak periods, this may not allow sufficient time. Also, it may not be practical in an enterprise which operates in multiple time zones.

We illustrate loading the IP network in Figure 13.

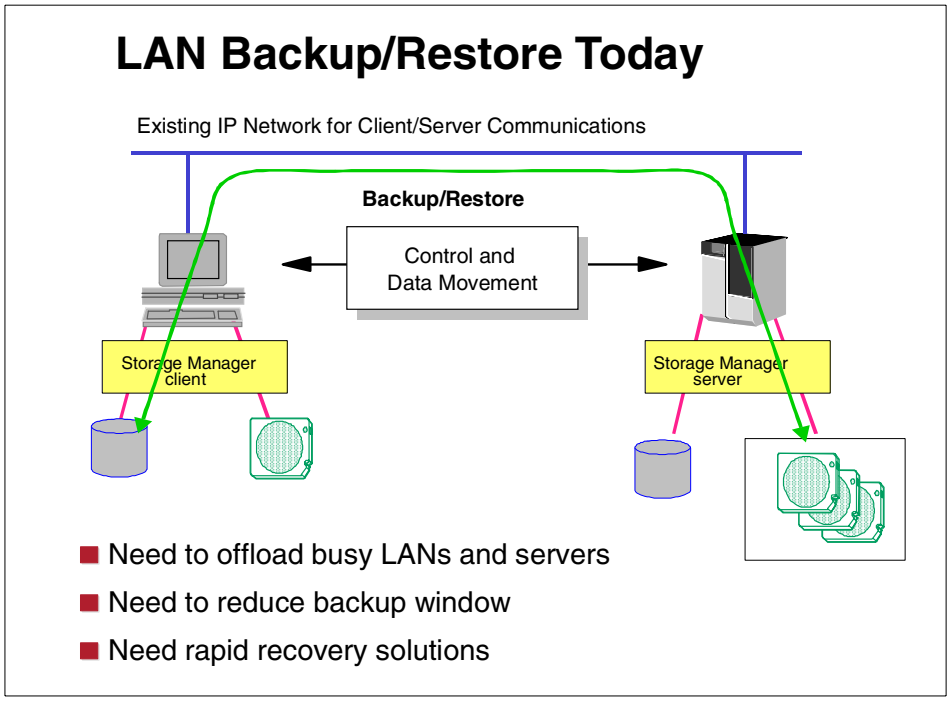

Figure 13. LAN backup/restore today - loading the IP network

SAN provides the solution, by enabling the elimination of backup and recovery data movement across the LAN. Fibre Channel's high bandwidth and multi-path switched fabric capabilities enable multiple servers to stream backup data concurrently to high speed tape drives. This frees the LAN for other application traffic. IBM's Tivoli software solution for LAN-free backup offers the capability for clients to move data directly to tape using the SAN. A

future enhancement to be provided by IBM Tivoli will allow data to be read directly from disk to tape (and tape to disk), bypassing the server. This solution is known as server-free backup.

1.5.5 High performance

Applications benefit from the more efficient transport mechanism of Fibre Channel. Currently, Fibre Channel transfers data at 100 MB/second, several times faster than typical SCSI capabilities, and many times faster than standard LAN data transfers. Future implementations of Fibre Channel at 200 MB/second and 400 MB/second have been defined, offering the promise of even greater performance benefits in the future. Indeed, prototypes of storage components which meet the two gigabit transport specification are already in existence, and may be in production in 2001.

The elimination of conflicts on LANs, by removing storage data transfers from the LAN to the SAN, may also significantly improve application performance on servers.

1.5.6 High availability server clustering

Reliable and continuous access to information is an essential prerequisite in any business. As applications have shifted from robust mainframes to the less reliable client/file server environment, so have server and software vendors developed high availability solutions to address the exposure. These are based on clusters of servers. A cluster is a group of independent computers managed as a single system for higher availability, easier manageability, and greater scalability. Server system components are interconnected using specialized cluster interconnects, or open clustering technologies, such as Fibre Channel - Virtual Interface mapping.

Complex software is required to manage the failover of any component of the hardware, the network, or the application. SCSI cabling tends to limit clusters to no more than two servers. A Fibre Channel SAN allows clusters to scale to 4, 8, 16, and even to 100 or more servers, as required, to provide very large shared data configurations, including redundant pathing, RAID protection, and so on. Storage can be shared, and can be easily switched from one server to another. Just as storage capacity can be scaled non-disruptively in a SAN, so can the number of servers in a cluster be increased or decreased dynamically, without impacting the storage environment.

1.5.7 Improved disaster tolerance

Advanced disk arrays, such as IBM's Enterprise Storage Server (ESS), provide sophisticated functions, like Peer-to-Peer Remote Copy services, to

address the need for secure and rapid recovery of data in the event of a disaster. Failures may be due to natural occurrences, such as fire, flood, or earthquake; or to human error. A SAN implementation allows multiple open servers to benefit from this type of disaster protection, and the servers may even be located some distance (up to 10 km) from the disk array which holds the primary copy of the data. The secondary site, holding the mirror image of the data, may be located up to a further 100 km from the primary site.

IBM has also announced Peer-to-Peer Copy capability for its Virtual Tape Server (VTS). This will allow VTS users to maintain local and remote copies of virtual tape volumes, improving data availability by eliminating all single points of failure.

1.5.8 Allow selection of “best of breed” storage

Internal storage, purchased as a feature of the associated server, is often relatively costly. A SAN implementation enables storage purchase decisions to be made independently of the server. Buyers are free to choose the best of breed solution to meet their performance, function, and cost needs. Large capacity external disk arrays may provide an extensive selection of advanced functions. For instance, the ESS includes cross platform functions, such as high performance RAID 5, Peer-to-Peer Remote Copy, Flash Copy, and functions specific to S/390, such as Parallel Access Volumes (PAV), Multiple Allegiance, and I/O Priority Queuing. This makes it an ideal SAN attached solution to consolidate enterprise data.

Client/server backup solutions often include attachment of low capacity tape drives, or small automated tape subsystems, to individual PCs and departmental servers. This introduces a significant administrative overhead as users, or departmental storage administrators, often have to control the backup and recovery processes manually. A SAN allows the alternative strategy of sharing fewer, highly reliable, powerful tape solutions, such as IBM’s Magstar family of drives and automated libraries, between multiple users and departments.

1.5.9 Ease of data migration

Data can be moved non-disruptively from one storage subsystem to another using a SAN, without server intervention. This may greatly ease the migration of data associated with the introduction of new technology, and the retirement of old devices.

1.5.10 Reduced total costs of ownership

Expenditure on storage today is estimated to be in the region of 50% of a typical IT hardware budget. Some industry analysts expect this to grow to as much as 75% by the end of the year 2002. IT managers are becoming increasingly focused on controlling these growing costs.

1.5.10.1 Consistent, centralized management

As we have shown, consolidation of storage can reduce wasteful fragmentation of storage attached to multiple servers. It also enables a single, consistent data and storage resource management solution to be implemented, such as IBM's StorWatch tools, combined with software such as Tivoli Storage Manager and Tivoli SAN Manager, which can reduce costs of software and human resources for storage management.

1.5.10.2 Reduced hardware costs

By moving data to SAN attached storage subsystems, the servers themselves may no longer need to be configured with native storage. In addition, the introduction of LAN-free and server-free data transfers largely eliminate the use of server cycles to manage housekeeping tasks, such as backup and recovery, and archive, and recall. The configuration of what might be termed "thin servers" therefore might be possible, and this can result in significant hardware cost savings to offset against the costs of installing the SAN fabric.

1.5.11 Storage resources match e-business enterprise needs

By eliminating islands of information, typical of the client/server model of computing, and introducing an integrated storage infrastructure, SAN solutions match the strategic needs of today's e-business. This is shown in Figure 14.

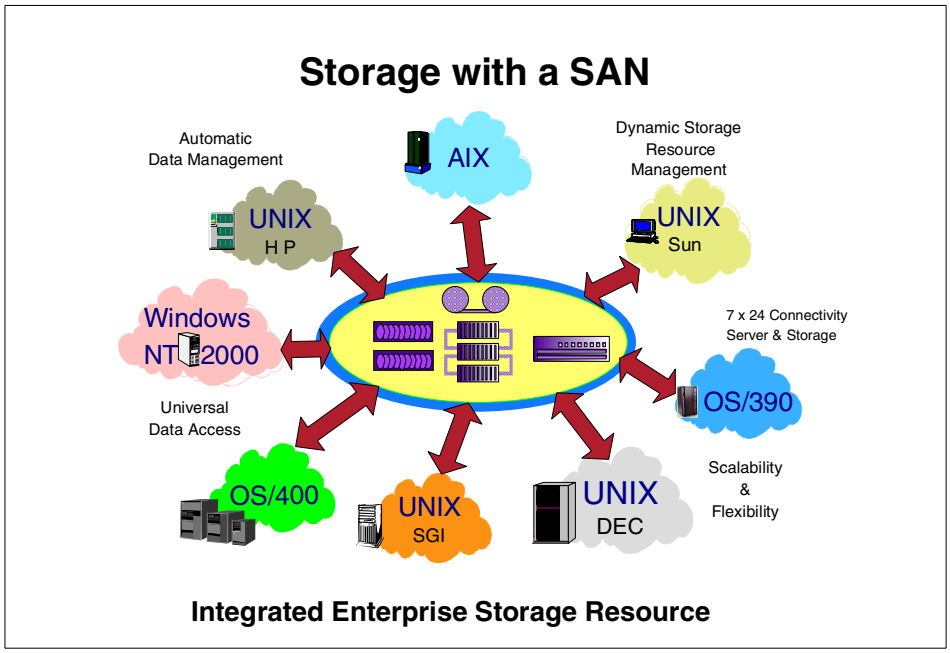

Figure 14. SAN solutions match e-business strategic needs

1.6 SAN market trends

In view of SAN's potential to deliver valuable business benefits, we should not be surprised at the substantial interest being shown by users, vendors and analysts alike. While early adopters have been installing limited SAN solutions since 1998, significant awareness among business users began to be generated during 1999. Many vendors announced SAN products and solutions in 1999, and this trend is accelerating in the year 2000. Analysts now estimate that industry revenue for network attached storage (both SAN and NAS), will grow rapidly during the next two years. Indeed, by the year 2003, external storage consulting groups estimate that SAN attached disk arrays will reach 48% of the revenue for externally attached disk arrays. NAS is expected to reach 23%, while disk arrays attached in the traditional manner directly to servers will account for only 29%. This is a dramatic shift in the IT infrastructure in a very short time frame.

It is also no surprise that the main drivers for SAN solutions are coming from the open systems environment. In 1999, for the first time, industry revenues for open systems disk arrays (UNIX, Windows NT, and so on) are estimated to have overtaken revenues for S/390 attached arrays. By the year 2003,

external storage consulting groups estimate that disk array revenues from the open environment will be approximately six times greater, while S/390 array revenues will remain relatively flat.

IBM's own estimates are, that by the year 2003, some 70% of all medium and large enterprises will install Fibre Channel SAN solutions to address varying business needs.

Stock markets around the world, especially Wall Street, are excited by the opportunities offered by the emerging Fibre Channel technology, and this is reflected in soaring stock prices of specialist manufacturers and developers of fabric components.

As with any new technology, it is up to the user to assess its relevance and value; and to decide if and when to deploy appropriate solutions. But buyers should also beware. It is easy to assume that, because other mature network solutions operate in a particular way (for example, in relation to the interoperability of solution components) so does Fibre Channel. This is not necessarily the case, because Fibre Channel standards for storage networking are still emerging. The purpose of this book is to discuss some of the factors which IT architects and planners should take into consideration, as they begin to investigate and to design business oriented SAN solutions for their enterprise.

Chapter 2. SAN planning and design considerations

Once you have concluded that Storage Area Networks will be beneficial to your organization, where do you go from here? How do you plan to go about implementing a SAN? Who should be involved? What do you need to do before you call your vendor and order new equipment? This chapter discusses some of the things you need to consider when you are at the planning and design stage. It does not purport to be a detailed methodology. It is intended only to cover some basic ideas and suggestions. IBM's International Global Services division offers detailed planning, design, and consultant services, which give a structured approach to SAN design and implementation.

2.1 Establishing the goals

There is an old maxim which states, "If it isn't broken, don't fix it". This can easily be applied to discussions about implementing SANs. When you look at your current storage infrastructure, if you find that it meets all your expectations for data availability, performance, and connectivity, then implementing a SAN will be difficult to cost justify. Most IT executives will be reluctant to make investments in a new IT infrastructure unless they can be shown that real benefits will accrue.

2.1.1 Business goals

As we have seen in 1.5, "The business benefits of a Fibre Channel SAN" on page 17, there are numerous ways in which a SAN can benefit an organization. Each company will have a unique set of circumstances and needs. If you can identify specific applications which today suffer from lack of connectivity; inability to share information or storage resources with others; which cannot be backed up in a timely manner due to bandwidth constraints on your LAN; or otherwise are limited in the way in which they provide service to the organization, then a SAN can be the solution. If users are always asking for more storage, and your storage costs are growing rapidly, and management of resources is becoming increasingly difficult, then a SAN is a likely answer. If your company is moving into e-business, supporting application operations 24 hours, 7 days a week, implementing Enterprise Resource Planning and Business Intelligence, and cannot tolerate outages of such mission critical applications, then a SAN can solve your problems.

In other words, you need to identify the "pain levels" associated with data movement, data sharing, data growth, and so on, in your own organization. Then, you can quantify how a SAN will contribute to your ability to achieve the

levels of service demanded by your business managers. This might be quantified in terms of improved communications within the business, and externally with customers. It can be a matter of improving the ability of managers and employees to make good business decisions due to better information availability. You might be measuring the security and reliability of mission critical applications, or reducing costs of storage hardware and skilled human management resources. The need may be to establish flexible, adaptable IT services, to stay in the race with competitors who are not constrained by legacy applications or inflexible IT infrastructures. The focus can be on the ability to increase revenue and profit with effective growth of new e-business services.

Whatever the level of “the pain”, you need to understand your existing IT infrastructure, which are the mission critical applications, and what are the business goals and directions.

2.1.2 Technical goals

When you understand the business goals, these will lead you to evaluate the technical requirements placed on the supporting IT infrastructure; and what a SAN must provide to meet these requirements. These can be measured in terms of reliability, availability, and serviceability (RAS); performance; scalability; security; manageability; and affordability.

2.1.2.1 RAS

Assess the nature of the applications to be supported on the SAN. Do they need to be available on a 24 hour, 7 days a week basis? If they are not mission critical, how much downtime, if any, is acceptable? What are the costs of downtime? Non-availability of some applications may be measured in hundreds of thousands or even millions of dollars per hour; for others there may be very limited financial impact. The answer may lead you to focus on hardware aspects, such as mean time between failure (MTBF), mean time to repair; and serviceability characteristics, such as fault tolerance, hot swappable components, failover facilities, error reporting, and call home capabilities.

2.1.2.2 Performance

What are the performance characteristics required to support the various applications on the SAN? How do you measure this? With throughput (MB/second) or I/Os per second, or response time? What is the maximum capacity or bandwidth required for peak loads? What percentage of the SAN capacity will be used on average, and at what level of utilization would it become saturated? What happens to performance in the event of failure of

SAN components? Can sufficient spare bandwidth be provided to continue to deliver acceptable performance?

2.1.2.3 Scalability

How much growth is expected? Will the SAN you design be required to support additional applications? If so, in what time scale, for instance within the next year or two years? How fast is data growing, and will you need to expand storage resources, or add more servers? Do you need to support legacy SCSI hardware? What are the distances between server and storage resources, and will this need to expand to include other departments and locations?

2.1.2.4 Security

How will you protect application data on the SAN from loss or corruption, without losing access to the information? Can you provide backup and recovery and disaster protection capabilities to meet your organizations policies for the data? What failover facilities within the SAN design will be required to ensure continued accessibility in the event of errors or disasters?

If multiple servers are attached to the SAN, can you ensure that each may only access the data or storage devices it is authorized to access? This is particularly critical in a heterogeneous platform environment, especially if Windows NT hosts are participating. For instance, Windows NT expects to see a SCSI bus attachment, and it seeks to attach all the devices which are attached to the bus. It does the same on a SAN, so it is essential to provide security against this occurrence by means of zoning and/or logical unit number (LUN) masking. Decide which level of zoning (hardware and/or software) and LUN masking is appropriate, remembering that LUN masking at the storage device level (including SAN Data Gateway) provides the highest level of security because it logically binds storage volumes to specific servers. This ensures that each server can only access its own data, just as though the storage was directly attached to the server.

2.1.2.5 Manageability

Consider how the SAN will be managed in your environment. You will need to have tools to handle a number of critical aspects. A variety of software vendors offer tools to address some of these requirements. Tivoli Systems has a set of software tools which provide complete SAN management.

These management tools include:

- **Configuration**: Facilities to identify, operate, collect data from and control the devices in the SAN

- **Access**: Tools to allow configuration, maintenance, zoning and LUN masking to protect data and ensure only authorized access to information
- **Performance**: Managing performance to meet service levels, analyze the traffic on the SAN, and understand the behavior of applications in order to be able to optimize the network and plan for future requirements
- **Faults**: The ability to detect, isolate, correct, and report on events within the SAN

2.1.2.6 Affordability

When you are considering a SAN configuration, you have already established the expected benefits. But costs are always one of the most significant aspects of any investment decision. What are the decision criteria which are most important in the solution you are contemplating? For a mission critical enterprise application it may be that high availability is the overriding requirement. In a campus-wide application, the dominating theme may be connectivity. At an individual departmental level, low cost may be the main objective. As with most other choices, you can make design trade-offs (Table 1), but each compromise usually involves giving up on something, whether it is performance, availability, security, scalability, manageability, or some other characteristic.

Table 1. Design trade-offs

Design Goal	Trade-off
High availability	Redundant components and higher costs
High performance	Higher cost circuits and more equipment
High level of security	More costly monitoring and reduced ease of use
High scalability	Higher costs with possible availability impacts
High throughput for one application	Lower throughput for another application
Low cost	Reduced availability and performance

2.2 Defining the infrastructure requirements

If you are starting from scratch with a totally new network in a green field site then you can go straight ahead with selection of the optimum SAN topology to meet your needs. But in most situations it is likely that you are replacing an existing infrastructure for storage. You may even be planning to change or upgrade an existing SAN implementation. So, before selecting a design for

the new SAN, it makes good sense to fully understand what it is that is being replaced. The current storage configuration, LAN or SAN network structure, application uses, traffic loads, peak periods and performance, as well as current constraints, are all relevant information in determining realistic goals for the SAN. This information will also help you to determine what, if any, of the existing components can be used in a new topology; and what will be involved in migrating from today's environment to the new one.

2.2.1 Use of existing fiber

In many cases you may already have fiber-optic cables laid in your organization. IT budget holders will want to know if you can use the existing cabling. This is discussed in greater depth in *Designing an IBM Storage Area Network*, SG24-5758. If the existing cabling has been laid for some time the answer may well be that the high speeds and accuracy required of Fibre Channel requires new cable investments. It is possible to test if installed fiber meets the necessary quality, but this can also be a costly exercise. If recent fiber cable has been laid you may need to decide what extensions need to be added to the configuration.

2.2.2 Application traffic characteristics

Before selecting a SAN topology you will need to understand the nature of the estimated traffic. Which servers and storage devices will generate data movements. Which are the sources, and which are the targets? Will data flow between servers as well as from servers to storage? If you plan to implement LAN-free or server-free data movement, what are the implications? How much data will flow directly from storage device to storage device, such as disk to tape, and tape to disk? What is the protocol? For instance, is this standard SCSI, or are you including digital video or audio?

What are the sizes of data objects sent by differing applications? Are there any overheads which are incurred by differing Fibre Channel frames? What Fibre Channel class of service needs to be applied to the various applications? Which departments or user groups generate the traffic? Where are they located, what applications does each community use, and how many in the user group? This information may point to opportunities for physical storage consolidation. It will also help you to calculate the number of Fibre Channel nodes required, the sum of all the data traffic which can be in transit at any time, and potential peaks and bottlenecks.

Can you identify any latent demand for applications, which are not carried out today because of constraints of the existing infrastructure? If you introduce high speed backup and recovery capabilities across a SAN, can this lead to

an increase in the frequency of backup activity by user groups? Perhaps today they are deterred by the slow speed of backups across the LAN? Could the current weekly backup cycle move to a daily cycle as a result of the improved service? If so, what would this do to SAN bandwidth requirements?

2.2.3 Platforms and storage

How many servers and what are the operating platforms which will be attached to the SAN? The majority of early SAN adopters have tended to implement homogeneous installations (that is, supporting a single operating platform type, such as all Netfinity, all HP, or all Sun servers). As SANs are maturing, the trend is towards larger scale networks, supporting multiple heterogeneous operating platforms (combining AIX, UNIX, Windows NT and so on) which has implications for security.

Fibre Channel capable servers require Fibre Channel HBAs to attach to the SAN fabric. The choice of HBA is probably already decided by the server vendor. Before you decide how many HBAs you require in your host to achieve optimal performance, you need to evaluate the performance of the server. Fibre Channel HBAs today transfer data at 100 MB/s. Can the system bus provide data at the same or higher speed? If not, the HBA will not be fully utilized. The most common system bus in use today is the Peripheral Component Interconnect bus (PCI), which operates at either 132 MB/s or 264 MB/s. Sun SBus operates at 50 MB/s, and HP HSC at only 40 MB/s. If the system bus delivers 132 MB/s or less, you will only need to attach one Fibre Channel HBA to the bus to achieve the required performance, since two would over run the bus speed. If you attach a second HBA it should only be for redundancy purposes. Our recommendation is to install one adapter per system bus.

Another major component of your current assets are the storage systems. You may have a variety of internally attached disk devices, which will not be relevant in a SAN operation. Also, you may have externally attached JBODs or RAID disk subsystems, and tape drives or libraries, which can be utilized within the SAN. These current assets have implications for the selection of interconnections to the SAN. You may want to support existing hardware which are SCSI or SSA compatible, and which will need to be provided with router or gateway connections for protocol conversion to Fibre Channel.

2.3 Selecting the topology

The most fundamental choice in the design of your SAN is the selection of the most appropriate topology. This selection may be colored by the overall approach to SAN planning that your organization wants to adopt.

The question is — top down, or bottom up design? In other words, should you try to design a corporate strategy, with a view to implement an enterprise wide SAN, or should you address the problem from the perspective of individual departments or user groups, and implement multiple SANlets? Maybe these small SANs will later merge into an enterprise-wide solution.

This is a difficult question to answer. Probably it will be answered differently depending on the size of the organization, the IT management philosophy, the politics of the organization, and the business objectives. It is also colored by the degree of risk which you associate with the implementation of an enterprise wide SAN today. The technologies are still relatively new. Industry standards in some key areas are still to be agreed upon. Not all server platforms can easily participate in Fibre Channel configurations, and the rate of change is extremely rapid. It is probable that in a years time things will look very different than they do today.

The fact is that the majority of SANs which have been implemented today are relatively small, point solutions. By this we mean that they were designed to address a specific “pain” or problem. Many users have implemented simple point to point Fibre Channel solutions to solve distance or performance issues. Many others have installed small clustered server solutions, or shared storage capacity by means of FC-ALs, because this provides improved connectivity and better utilization of storage resources. Others have designed switched fabric solutions for departments or have used FC directors to facilitate large scale storage consolidation in campus locations.

In practice then, the bottom up approach seems to be the most pragmatic. Solve specific application needs now to deliver value to your organization. This does not mean that you should not establish some common guidelines or standards regarding the purchase of equipment within the enterprise. This will facilitate interoperability in the future and avoid dead end investments, which cannot be integrated in a larger SAN environment as you expand the topology in the future. You may decide that there are a number of discrete and independent operating environments within your organization, and these will not need to be interlinked in the future. If so, you may choose to establish SAN islands which are configured with different topologies and components in which cross island interoperability is not required.

A strong trend is towards switched fabric environments. This is because fabric offers greatest flexibility and scope for the future. You may choose to install small FC-AL topologies now, for reasons of cost, or because the application being addressed is small scale today. If so, there is good logic in selecting hardware such as the IBM 3534 Managed Hub. This gives you flexibility for the future, such as more sophisticated function, manageability, and upgrade-ability, and with compatibility within a family of fabric devices.

Remember that FC-AL is not designed for high performance. It does not scale. As you add more devices on a loop, performance will tend to reduce because of the shared bandwidth and arbitration overheads. A maximum of two servers is advisable on a FC-AL. If one server fails and has to reboot, causing a new Loop Initialization Primitive sequence (LIP), it will bring down the whole loop. Availability is, therefore, also a serious consideration.

FC-SW, on the other hand, scales performance as you add nodes. This does not apply to inter switch links (ISLs), which do not add bandwidth between end nodes, because ISLs reduce the number of end to end connections. Secure zones and masked LUNs can be created so that only authorized servers can access specific information. FC-SW provides comprehensive, flexible growth options, but is more expensive at the outset.

2.3.1 Assessing the components

IBM provides a hierarchy of interconnect options with which to build Fibre Channel SANs to suit differing application characteristics. These can range from FC-AL topologies for work groups, to departmental switched fabrics, and to highly available topologies based on cascaded switches or fault tolerant directors. In addition, bridge solutions allow for attachment of legacy SCSI and SSA devices. This hierarchy is illustrated in Figure 15.

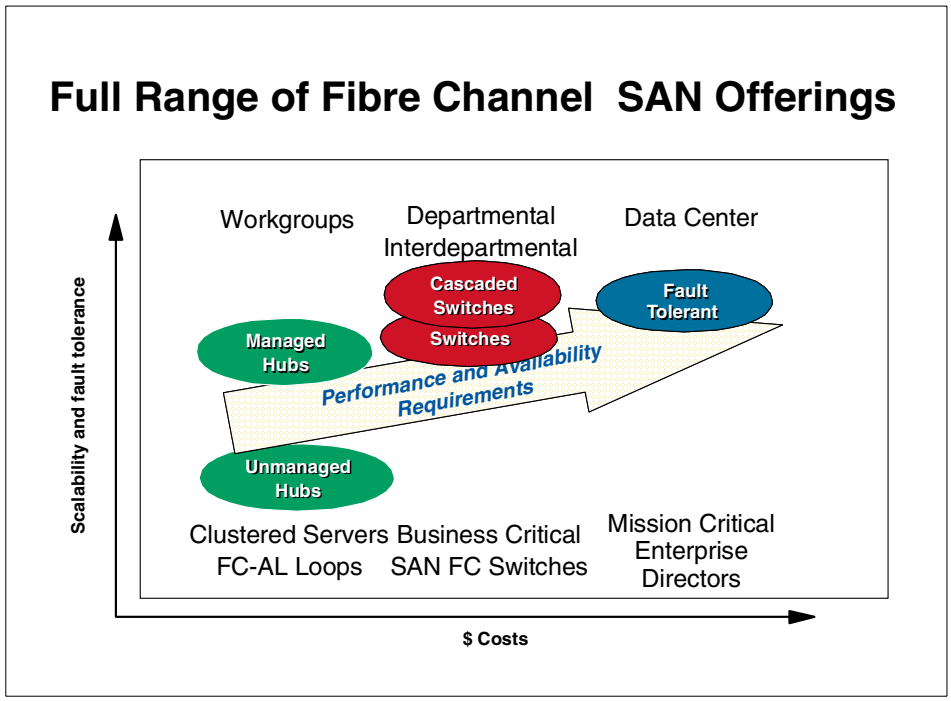

Figure 15. IBM's hierarchy of Fibre Channel SAN offerings

It is worth bearing in mind that new products, features and functions are regularly announced. For the latest information on the products described in this book, and for details of new solutions, refer to IBM's SAN Web site:

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

The following section summarizes the major reasons for selecting a device type.

2.3.1.1 When will you use a Router or a Data Gateway?

Routers and Gateways act as "bridges" between different protocols. You would select them to allow you to provide investment protection for IBM or non-IBM storage devices which use SCSI or SSA protocols, and provide attachment to the Fibre Channel SAN.

The IBM SAN Data Gateway Router (IBM 2108-R03) is a low-cost solution supporting attachment between a Fibre Channel attached host and a SCSI tape library, such as a 3575 Magstar MP.

The Vicom Fibre Channel SLIC Router Model FC-SL (7139 model 111) enables all IBM 7133, 7131, and 3527 SSA Disk Systems to attach to host systems using fibre channel host adapters and drivers.

The IBM SAN Data Gateway (2108-G07) provides protocol conversion for connection of SCSI and Ultra SCSI storage devices to Fibre Channel environments using an industry standard Fibre Channel Arbitrated Loop (FC-AL) interface. The SAN Data Gateway enables SCSI devices to benefit from distance extension to 500 meters, increased bandwidth of Fibre Channel, and increased addressability.

A wide range of IBM and non-IBM SCSI based servers are supported (including UNIX and Windows based), plus many IBM and non-IBM SCSI storage devices, including IBM Magstar tape and the Enterprise Storage Server. Because of its comprehensive zoning access control capabilities, including persistent binding of hosts to LUNs (LUN masking), the SAN Data Gateway is an ideal solution to support attachment of a SCSI based ESS to multiple hosts in a Fibre Channel SAN for storage consolidation.

2.3.1.2 When will you use a hub?

You will use a hub to implement a Fibre Channel Arbitrated Loop. Hubs can also be used as distance extenders, in connection with the IBM SAN Data Gateway.

Usually they are used for entry level homogeneous server implementations. Some of the possible uses of these hubs are clustering, LAN-free backup, storage consolidation, and remote disk mirroring.

The IBM products available are the IBM Fibre Channel Storage Hub (2103-H07) and IBM Fibre Channel Managed Hub (3534-1RU). The IBM 3534-1RU offers superior function due to its manageability, with superior fault isolation, planning, and controlling. It also has a non-blocking architecture. This means that any two pairs of ports can be active and transferring data, without blocking the transfer of data from another pair of ports, therefore, guaranteeing full-speed data delivery irrespective of traffic conditions. This product, technically, is based on the IBM 2109 SAN Fibre Channel Switch. Therefore, it has the potential to be made upgradable in the future, so possibly protecting your investment. For these reasons, we recommend that you normally select the IBM 3534-1RU in preference to the unmanaged hub.

2.3.1.3 When will you use a switch?

You will use a switch as the basis for development of a full switched fabric SAN. IBM's products today are the IBM SAN Fibre Channel Switch, with two models, an 8-port model and a 16-port model (IBM 2109 S08 and S16).

Multiple switches can be interlinked (cascaded) to build a large SAN comprising many ports. The ultimate limitation in the fabric design is 239 physical switches (imposed by the maximum number of unique domain IDs that can be defined). Today the practical tested limit is about 15% of this number, and with no more than seven hops allowed from the source port to the destination port. There are fabric designs in production today with between 10 and 20 switches in a single fabric. This number will certainly grow significantly over time.

The IBM 2109 also supports attachment of FL-Ports, so it can interlink to Fibre Channel Arbitrated Loops. In addition, hosts which are not fabric aware, and only operate on a Private Loop, such as HP servers, can be supported on the IBM 2109 using a feature known as QuickLoop (QL).

The switch can be set up to create a logical private loop with the storage assigned to that server. The whole switch can operate in QL mode, or individual ports can be configured as QL. In this way the IBM 2109 can be used instead of a hub to support such servers. The IBM 2109 can also be cascaded with the IBM 3534-1RU Fibre Channel Managed Hub. Therefore, the switch can be used in a number of ways for cost, availability and performance, to satisfy differing SAN application and user group needs within the enterprise.

2.3.1.4 When will you use a director?

You would select the McDATA ED-5000 Fibre Channel Director for high availability applications requiring extensive fault tolerance within the switch, and high port count to support multiple node attachments and high switch bandwidth. These applications might typically be found in large data centers, supporting large numbers of heterogeneous open systems servers.

The McDATA ED-5000 is based on the design of the IBM ESCON Director, which is widely used in S/390 Data Centers for core mission critical applications. The McDATA ED-5000 Director will support cascading of directors using E-Ports, but it does not provide FL-Port connectivity to arbitrated loops. Also, due to differences in implementation of name server and zoning techniques, the director is incompatible with the IBM 2109 Fibre Channel Switch and IBM 3534 Managed Hub, so they cannot yet be used together in a cascaded fabric. This may be resolved in the future as standards are agreed and implemented. The McDATA ED-5000 Director does support servers and devices attached using the IBM SAN Data Gateway and the IBM SAN Data Gateway Router.

2.3.2 Building a multiswitch fabric

A single switch or director is limited in the number of ports it can directly interconnect. To increase connectivity in the fabric it is necessary to connect multiple switches or directors. This is known as a cascaded fabric.

2.3.2.1 Cascading

Cascaded fabric is a cost effective, reliable way to achieve very large port counts in the SAN. It is also used as a means of delivering fault tolerant fabrics by eliminating single points of failure, and to service applications requiring high availability. Cascading also increases the maximum distance between interconnected devices. Examples used in this chapter will be based on Inter Switch Links (ISLs) between multiple IBM 2109 Fibre Channel Switches, (the first cascadable Fibre Channel switch available on the market).

2.3.2.2 Inter switch links

When cascading IBM 2109 switches the ports used for ISL will automatically be designated as E-Ports by the switch software. E-Ports reduce the number of ports available for device connection. More switches can be added to the fabric non disruptively. Multiple links can operate concurrently between any two switches in the fabric, allowing multiple redundant paths to be defined. All ISLs carry traffic. In the event of a link failure, the traffic it was carrying will be automatically and immediately transferred to other links. Adding ISLs will automatically cause routing and zoning information to be updated across all ISLs. Changing ISL configurations causes recalculation of routes within the fabric. This task is a load on all switches in the fabric, so numerous changes should be avoided. The maximum number of ISLs from one switch to a single adjacent switch is eight, but more than eight ISLs can be configured from a single switch if they attach to several other switches. This is shown in Figure 16.

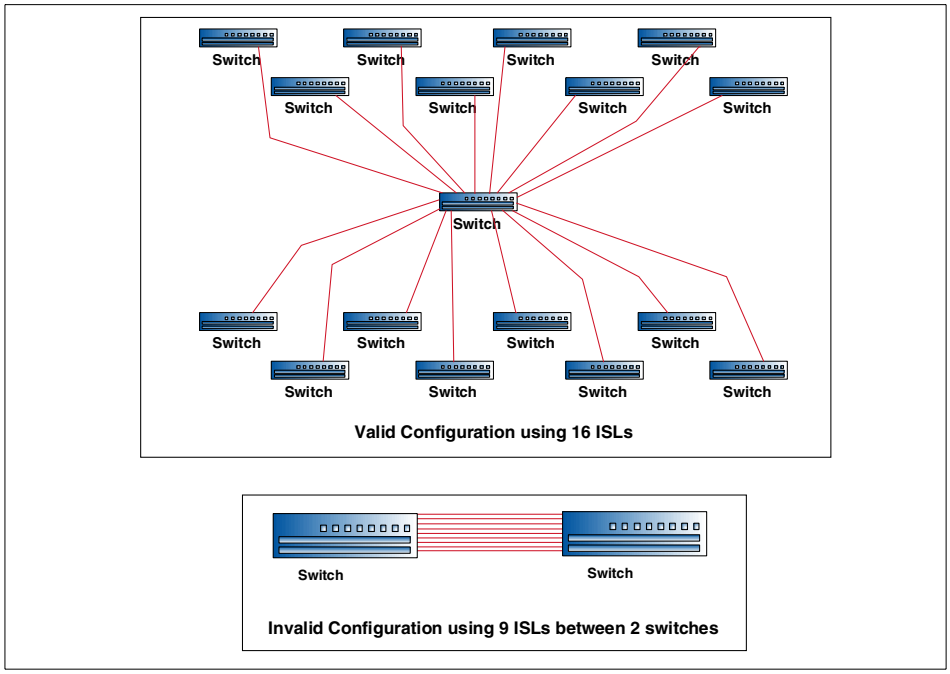

Figure 16. Valid and invalid inter switch links

2.3.2.3 Distributed name server

If a switch fails, the other switches in the fabric, and the nodes attached to them, are unaffected. Nodes attached to the failed switch are, of course, unable to talk to each other or to nodes on other switches. However, this problem can be overcome, since any node can have several Fibre Channel interfaces, each one attached to nodes on different switches in the fabric. This is illustrated in Figure 17. If any link fails, every switch can still communicate with all the other switches. IBM 2109 switches use a distributed fabric-wide name server. This means that the name server is fully distributed to each switch, therefore ensuring no single point of failure. When end nodes attached to servers and devices want to communicate to other nodes across the fabric, any switch provides information about the devices connected to the fabric by means of the distributed name server, even in the event of a failed switch.

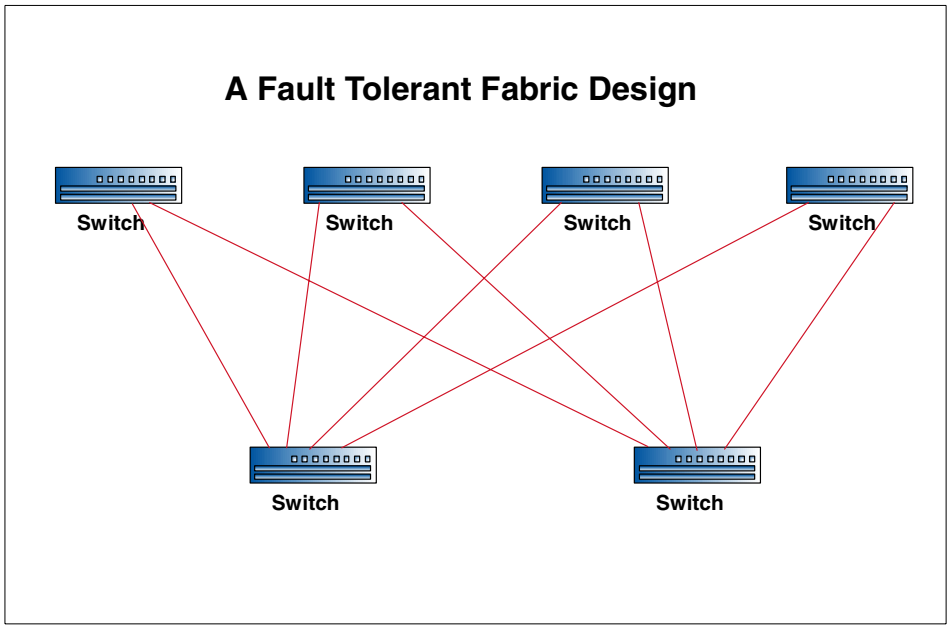

Figure 17. A fault tolerant fabric design

2.3.2.4 Fabric Shortest Path First (FSPF)

FSPF is the path selection protocol used by the IBM 2109 switch. It automatically calculates the best path between any two switches in the fabric when switches are powered up. It establishes all the routes across the fabric, and these change only in the event of a failure, or if a new ISL is created which offers an equal or better path to a given target. FSPF is very resilient to failures of hardware and software, automatically computing an alternate path around a failed link, typically in less than one second. If several equivalent paths are available between two nodes FSPF will automatically share traffic between these paths. This feature provides high bandwidth as well as fault tolerance, because no paths are held idle as stand-by redundant links. This is quite different to LAN path redundancy, which maintains idle paths for redundancy. FSPF can guarantee in-sequence delivery of frames, even if the routing topology has changed during a failure, by enforcing a “hold down” period before a new path is activated. This allows all frames in transit, to a specific destination, to be delivered or discarded.

An example of shared traffic routing is shown in Figure 18. Four switches have two ISLs between any two of them. Traffic between Switches 1 and 3, for example, will be shared on two paths, but traffic between Switches 1 and 4 can be shared on four paths.

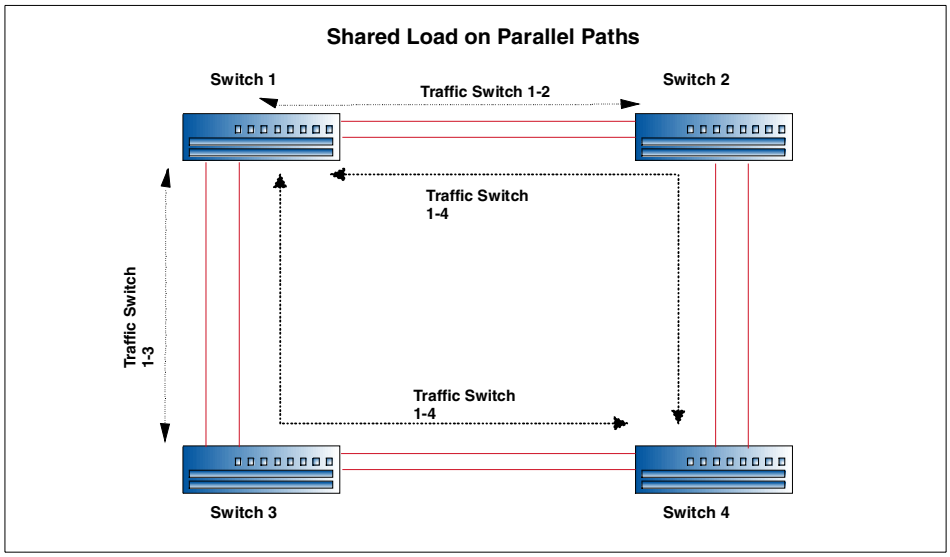

Figure 18. Load sharing on parallel paths

A maximum of seven hops is recommended between any two switches in the fabric to avoid time-outs. We say recommended, because the actual hops are not monitored and restricted to seven. More hops are possible and test beds up to 20 switches have been installed and tested, but every extra hop adds about 1.2 microseconds latency to the transmission. The length of the fiber is another consideration, since each kilometer between nodes adds a further five microseconds delay. Traffic patterns need to be understood, to avoid long paths and bottlenecks. Ideally devices should be attached to the same switch if they exchange large amounts of data, because this minimizes communication delays. If this is not possible then more ISLs should be configured to increase the available bandwidth between switches. Of course, this also adds to the resiliency of the fabric.

The fabric design in Figure 19 illustrates a fully meshed fabric in which a switch is only one hop from any other switch. This minimizes latency across the fabric. Also, if any link fails (even with two link failures) all switches can still communicate with each other.

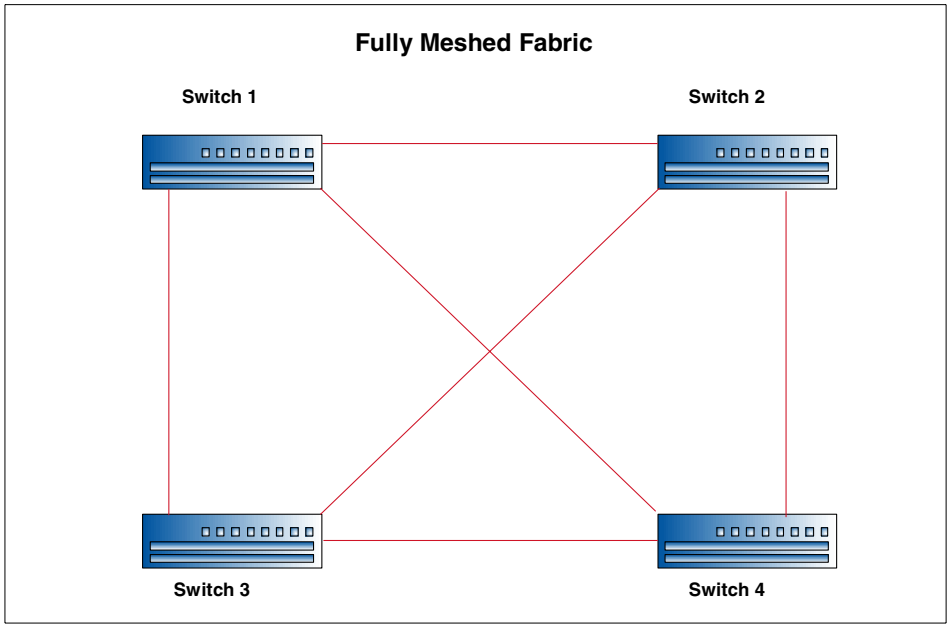

Figure 19. A fully meshed topology

2.3.2.5 Redundant fabrics

We have been discussing redundant elements, like paths and switches, within a single fabric. Another approach, and one which gives many advantages, is to use redundant fabrics. The simplest version of this is two switches which are not interconnected. If one switch fails, data is automatically routed via the second switch. This is initiated by host/device driver software, like IBM's Subsystem Device Driver (formerly Data Path Optimizer), which recognizes the failure of the path and fails over to the alternate path on the redundant switch. This configuration, illustrated in Figure 20, also allows for maintenance or repair actions to be made on one SAN, while the other stays in operation. More detailed examples of redundant fabric designs are described in *Designing an IBM Storage Area Network*, SG24-5758.

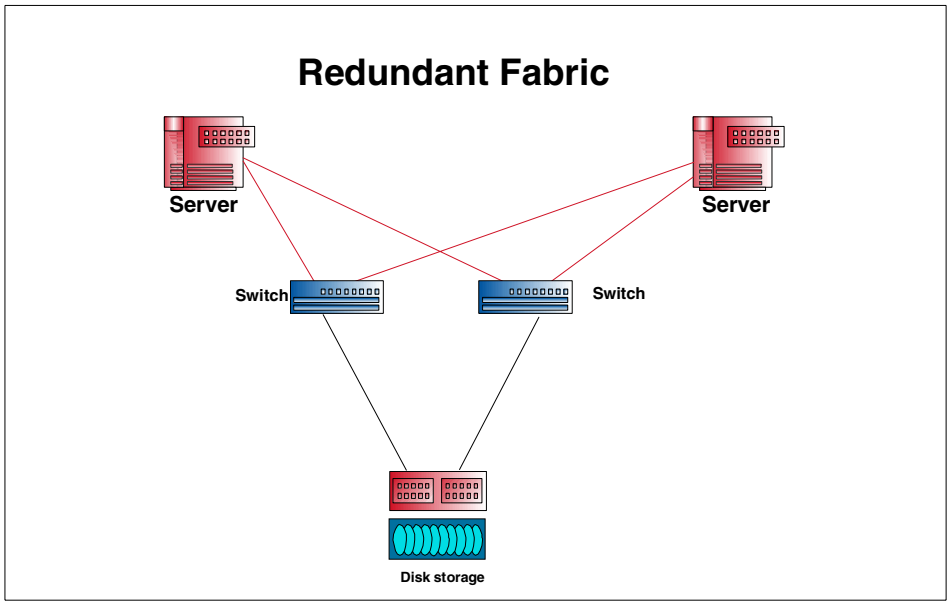

Figure 20. Redundant fabrics

2.3.2.6 Fabric backbone

Building on the concept of departmental SAN islands, each designed with fault tolerant fabrics, it is possible to link such islands together into an enterprise SAN (SAN continent perhaps), by providing a fault tolerant backbone of inter linked switches. This concept is shown in Figure 21. The fabric design shown here provides a total of 186 nodes in total, with up to 150 nodes at the department level, and 36 nodes in the backbone fabric. The backbone SAN can be used for shared devices such as a tape library, which can be accessed from any node in the enterprise fabric.

In this manner you can begin to use building block SANlets, and grow these to larger, resilient cascaded fabrics with multiple alternate paths through the network. Such designs can be easily expanded using tiers of switches for redundancy and increased bandwidth. In such an environment you will need to use zoning and LUN masking to provide appropriate security to application data.

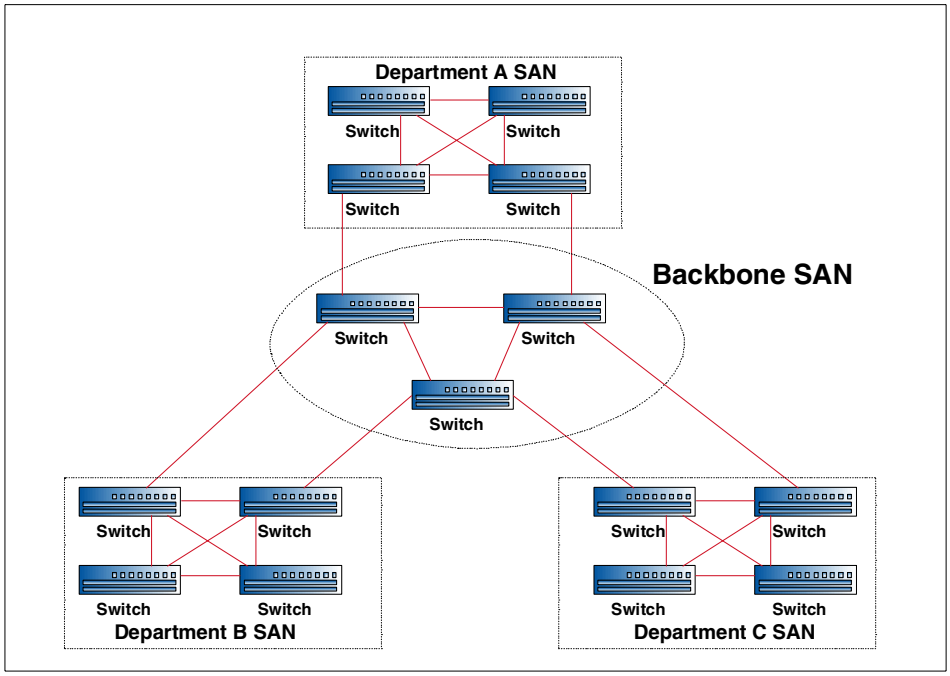

Figure 21. Fabric backbone interconnects SAN islands

2.3.3 Quality of service requirements

An important criterion for selection of SAN components relates to the level of service required from the SAN. This includes all aspects of the technology (hub, switch or director), the topology (loop or fabric), and the degree of redundancy, including fault tolerance. This is particularly relevant for organizations serving the global marketplace 24 hours per day, seven days per week over the Internet. In the e-business economy of today, continuous availability is not optional. If you are not online, you are not open for business, and widely reported incidents of system outages in well known e-business companies show that loss of revenue can be immense.

Strategic Research Corporation (SRC) has described the principle of Quality of Service (QoS). This is a framework used to establish appropriate performance and availability characteristics of a complex service, such as a network. Specifically, SRC defines service levels for SANs which they call Quality of Connection (QoC). This is built on the concept of system availability commonly used when defining service levels. These are normally described in terms of percentage systems availability.

A 99.999% (five 9s) up time refers to achievement of less than five minutes systems downtime in one year. A one 9 measure refers to a 90% availability (less than 36.5 days systems downtime), and a three 9s level is 99.9% uptime (less than 8 hours 45 minutes systems downtime annually). Downtime can be defined as any complete interruption of service for any reason, whether planned or unplanned.

To meet the very high levels of uptime required by planners and administrators it is essential to design the correct network architecture. It needs built in fault tolerance, failover capabilities, and available bandwidth to handle unplanned outages in a transparent manner. SAN QoC measurements are determined by the network topology, and the interconnect technologies used (hubs, switches). These define how well the SAN can sustain operations in the event of an outage within the network, from the perspective both of connection availability and maintaining performance.

High availability can be built in to the fabric by eliminating single points of failure. This is achieved by deploying hardware components in redundant pairs, and configuring redundant paths. Redundant paths will be routed through different switches to provide availability of connection. In the event of a path failure (for instance due to HBA, port card, fiber-optic cable, or storage adapter), software running in the host servers initiates failover to a secondary path. If the path failover malfunctions the application will fail. Then the only choice is to repair the failed path, or replace the failed device. Both these actions potentially lead to outages of other applications on multiple heterogeneous servers if the device affected is the switch.

Switches, like the IBM 2109, have redundant, hot pluggable components (including fans, power supplies, ASICs and GBICs), which can be replaced during normal operation. These hardware failures cause little or no noticeable loss of service. However, in the case of some failed components (such as the mother board) the switch itself will be treated as the field replaceable unit (FRU). Then all the ports and data paths are taken down. Automatic path failover will occur to another switch, so the network continues to operate, but in degraded mode.

Here there is a distinction between a switch and a director. Using the analogy of disk arrays, an individual switch can be likened to a JBOD in that it is just a bunch of ports. That is to say, although it has redundant components, in the event of certain component failures the total switch can fail, or must be replaced as the FRU. Cascading of multiple switches can achieve a higher level of fault tolerance. A single director can be viewed more like a RAID subsystem, in that it is designed to be highly fault tolerant. Only the failure of the mother board would result in total failure of the director. All other

components are redundant, with automatic failover. Redundant field replaceable units are hot swappable, and microcode updates can be made non disruptively. Maintenance capabilities, such as *call home* are supported.

According to tests run by CLAM Associates, it can take more than an hour to replace and reconfigure a new switch and bring it back into operation. For a 16-port switch this equates to 960 path minutes of degraded performance, as defined by the SRC QoC methodology. A path minute describes one user port being unavailable for one minute. Using path minutes as a way of describing the impact of an outage SRC defines five levels of QoC as shown in Figure 22.

SAN Quality of Connection

QoC Class	Fault Tolerance	Availability Annual uptime	Performance Degradation Path minutes/ year	Bandwidth Scaleability	Device Topology
1	Failure sensitive No redundancy	90%	Not applicable	Single Point to Point or Loop	Single Hub
2	Failure resilient Partially redundant paths and interconnects	99%	50,000	Loops and/or Switched Fabric	Single or Dual Hubs or Fabric Switches with single or dual paths
3	Failure resilient Fully redundant paths Fully redundant or fault tolerant interconnects	99.9%	5000	Switched Fabric	Dual Fabric Switches or Single Director
4	Failure tolerant Fully redundant paths and interconnects Fault tolerant backbone interconnects	99.99%	500	100% Switched Fabric + Max # of Ports per backplane	Dual Directors
5	Fault tolerant Fully redundant paths and interconnects All interconnects fault tolerant	99.999%	50	100% Switched Fabric + Max # of Ports per backplane	Dual Directors

Source: Strategic Research Corp.

Figure 22. SAN Quality of Connection

For instance, a single point to point topology has no redundancy and is classified as Class 1 QoC. If a failure occurs there is no access, so there is also 100% performance degradation. Class 2 has some redundancy with multiple paths and interconnects, but an outage can still occur for an extended period. Dual switches, or a single director provide full path and interconnect redundancy in Class 3, but a failure would imply degraded performance delivering variable service levels.

SRC defines Class 4 QoC as Failure Tolerant and Class 5 as Fault Tolerant. The Class 4 network must be able to recover from an outage and not incur more than 500 path minutes of degraded operation per year; and Class 5 must meet the five 9s measure of availability with only 5 minutes down time and 50 path minutes degradation. This requires multiple fabric connections between all devices, requiring director level hardware, or multiple inter switch links (ISLs) in a meshed fabric. A consideration with ISLs, using E-port connections, is that they do not add bandwidth to the total configuration, only bandwidth and connectivity between switches. For this reason, SRC concludes that to meet Class 4 and Class 5 QoC requirements today, for fault tolerance with scalable performance, the maximum number of ports per backplane are required; hence directors are favored for these mission critical applications. This is due to their larger number of ports (32), and n+1 fault tolerant design. Future switches with larger port counts would also address this effectively if configured in redundant meshed fabrics.

2.3.4 Hierarchical design

What we have seen is that a SAN can take numerous shapes. When you start thinking about SAN design for your own organization you can learn from the experience gained in the design of other, mature networks such as LANs, and the Internet. In these a hierarchical network structure has generally been adopted, to facilitate change, allow easy replication as the structure grows, and minimize costs. This hierarchy comprises three layers (as shown in Figure 23).

The core

At the center is a high speed, fault tolerant backbone, which is designed to provide very high reliability (QoC Class 4 or 5 as defined by SRC). This is designed to minimize latency within the fabric and to optimize performance. This core would normally be built around fault tolerant Fibre Channel directors, like the McDATA ED-5000, or a fully redundant, meshed topology of switches, like the IBM 2109.

The distribution layer

The distribution layer of the hierarchy would comprise fault resistant fabric components, designed to deliver QoC Class 3 or Class 4, depending on the applications. Good connectivity and performance would be prime considerations.

The access layer

Here are the entry point nodes to the fabric, comprising host bus adapters, routers, gateways, hubs, and switches that are appropriate to service the number of servers and storage devices supported on the fabric.

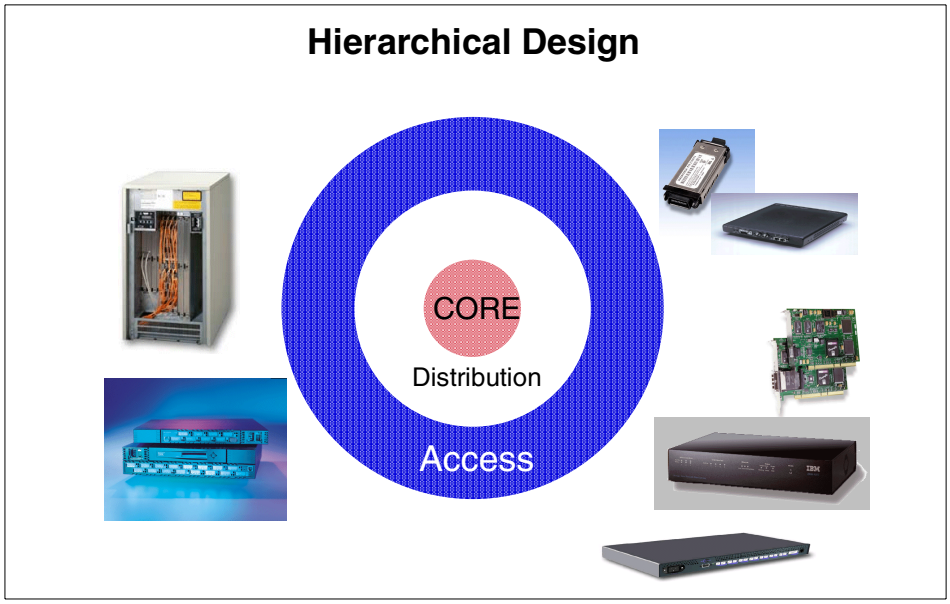

Figure 23. SAN hierarchical design

This hierarchy is analogous to the telephone switching system. Each user has access to the network using an individual node (the telephone); these link to local area switches, which in turn link to central core switches which serve a large national and international network with very high bandwidth. A similar hierarchy has been built to serve the Internet, with end users linked to local Web servers, which in turn communicate with large scale, high performance, core servers.

2.4 The next steps

Now that you are ready to design your SAN, there are many things to do.

2.4.1 The planning team

You will need to bring together the people with the appropriate skills to plan and implement the SAN. Who should be in the team depends on the scale of the project. This might range from installing a simple point to point connection to solve a distance issue in one location, to a large scale SAN comprising multiple meshed fabrics, interconnected across a large campus, or linking between several locations in order to serve a large organization. In the first case, the “team” may just be one or two storage administrators.

In the enterprise-wide case, you will probably need to include a number of skills. Since you are planning a network, it makes only common sense to include staff who have knowledge of complex networks. You will want to have representatives who know about the various platforms (UNIX, Windows NT, Novell Netware, AS/400 systems, Numa-Q, and so on), since there are differing system requirements and quirks which must be understood. Also consider the databases and applications you will be supporting, and include advisors for these. Knowledge of fiber cabling and data center planning may be necessary. You will certainly, of course, need strong storage planning and management skills that are appropriate to the platforms, subsystems, and software tools included in the design. Additionally, you will need project management skills to coordinate the whole exercise.

2.4.2 Equipment selection

The detailed list of logical and physical connections required in the SAN should act as the basis for defining your fabric hardware requirements and for arriving at an estimated implementation cost. Now you are ready to make the final selection with your vendor.

2.4.3 Interoperability testing

No doubt you will want to ensure that the SAN solution you are designing will operate correctly in practice. As industry standards are still under development this is particularly pertinent. You may select a pretested and certified solution, in which case there is little or no risk. Vendors throughout the industry are testing their hardware and software in many differing environments. IBM Global Services has made major investments in laboratories in the USA and Europe to help you with such testing. Details of tested and certified solutions are constantly being updated, and are posted on the IBM SAN Web site.

2.4.4 Documentation

As with any project, you will need to fully document the SAN plan. This should include details about most of the topics already discussed in this chapter:

- Business goals
- Technical requirements
- Infrastructure — current and planned
- Cabling
- Traffic characteristics today and expected
- Platforms and storage devices — current and planned
- SAN applications
- Logical and physical design

- Hardware and software plan
- Training plan
- Project implementation plan

2.5 Future developments

We can all speculate on what the future will bring for SANs. Fibre Channel version 2 will bring speeds of 200 MB/s, double what we have today. SAN fabric vendors will most likely develop new, more intelligent and faster gateways, hubs, and switches, with more scaleability in port count, bandwidth, and greater fault tolerance. Server and storage vendors will introduce more native Fibre Channel solutions; faster processors, more scalable and intelligent storage subsystems. Fibre Channel industry standards will continue to be delivered through cooperation among the vendors, creating greater ease of interoperability and rapid growth in the SAN marketplace. S/390 FICON and FCP protocols will be enabled to operate on the same fabric. Sophisticated software management tools will finally deliver full end-to-end SAN management. The true Enterprise SAN will arrive, as we show in Figure 24.

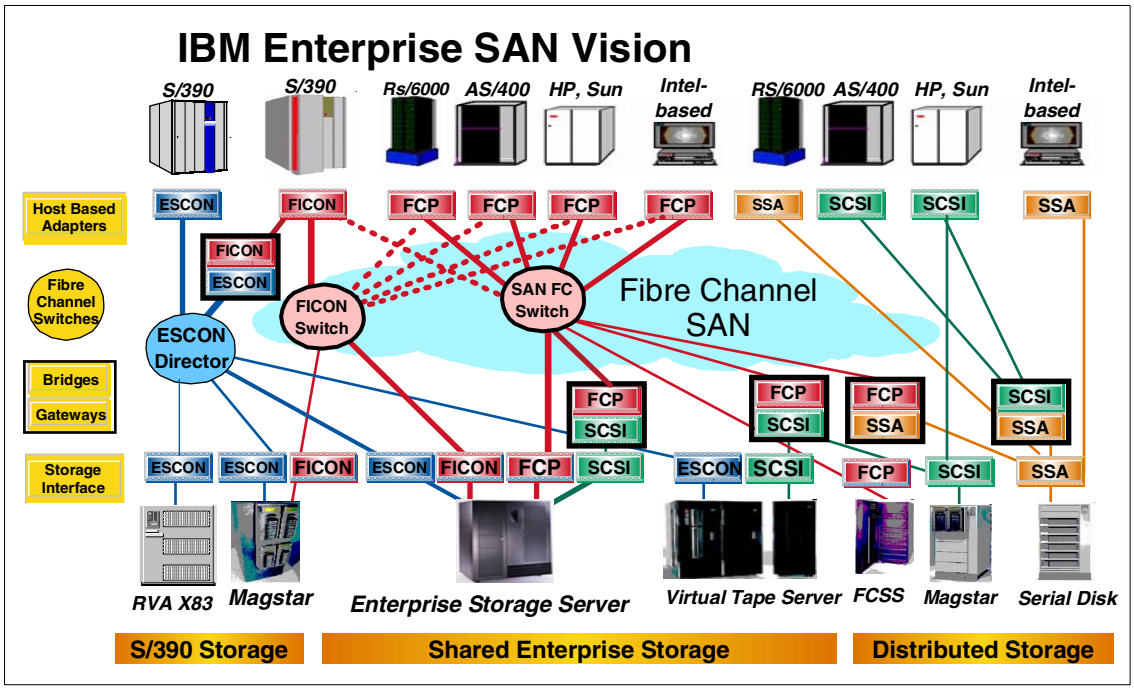

Figure 24. The IBM Enterprise SAN vision

In a fast moving world there is often a tendency to hold back and wait for something better which may be just around the corner. But those who wait are sometimes left behind in the race for competitive advantage. The race is on. Now is the time to join in.

Chapter 3. Implementing Fibre Channel host adapter cards

In this chapter we will describe the steps involved in implementing Fibre Channel host adapter cards in a supported environment.

This chapter lists the steps necessary to implement Qlogic Fibre Channel cards in Windows NT and RS/6000 host systems:

- 3.1, “Installing QLogic host adapter cards on a Windows NT system” on page 53
- 3.2, “Installing Fibre Channel host adapter cards on a RS/6000 system” on page 58

For more general information or technical information on employment and application of these adapters, refer to the appropriate Fibre Channel adapter manufacturer.

3.1 Installing QLogic host adapter cards on a Windows NT system

In the following sections, we cover these topics:

- 3.1.1, “Installing the QLogic 2100F adapter card” on page 53
- 3.1.2, “Installing the QLogic 2200F adapter card” on page 54
- 3.1.4, “Installing the drivers” on page 57
- 3.1.5, “Locating the World Wide Port Name (WWPN) on Windows NT” on page 57

3.1.1 Installing the QLogic 2100F adapter card

The Qlogic QLA2100F adapter card only supports FC-AL. This section lists the steps you have to follow to attach an ESS to a Windows NT host system with the Qlogic QLA2100F adapter card.

1. Install the QLogic QLA2100F adapter card in the host system.
2. Connect the cable to the ESS port.
3. Restart the server.
4. Press and hold the Alt key and press the Q key to get to the FAST!Util command panel.
5. From the Configuration Settings menu select Host Adapter settings.
6. From the Advanced Adapter Settings menu, use the down arrow key on your keyboard to highlight LUNs per target; then press Enter.
7. Use the down arrow key on the keyboard to highlight LUNs per target.

8. Use the down arrow key on the keyboard to find and highlight 256. Press Enter.
9. Press the escape key.
10. Save the changes. Highlight Yes.
11. Restart the server.

3.1.2 Installing the QLogic 2200F adapter card

This section tells you how to attach an ESS to a Windows NT host system with the Qlogic QLA220F adapter card. Perform the following steps to install the QLogic QLA2200F adapter card:

1. Install the QLogic QLA2200F adapter card in the host system.
2. Connect the cable to the ESS port.
3. Restart the server.
4. Press and hold the Alt key and press the Q key to get to the FAST!Util command panel.
5. From the Configuration Settings menu select Host Adapter settings.

   To set the parameters and values from the Host Adapter settings menu, use the following:

   a. Host adapter BIOS: Disabled
   b. Frame size: 2,048
   c. Loop reset delay: 5 (minimum)
   d. Adapter hard loop ID: Disabled

      Some configurations might require hard IDs. Consult the separate application note that addresses Fibre Channel address for more information, or consult the adapter card vendor.

   e. Hard loop ID: ?
6. From the Advanced Adapter Settings menu, use the down arrow key on your keyboard to highlight LUNs per target; then press Enter. Set the parameters and values from the Advanced adapter settings menu as follows:

   a. Execution throttle: 240
   b. Fast command posting: Enabled
   c. >4 GB addressing: Disabled for 32 bit systems
   d. LUNs per target: 0

e. Enable LIP reset: No

f. Enable LIP full login: No

g. Enable target reset: Yes

h. Login retry count: 20 (minimum)

i. Port down retry count: 20 (minimum)

j. Driver load RISC code: Enabled

k. Enable database updates: No

l. IOCB allocation: 256

m. Extended error logging: Disabled (might be enabled for debugging).

Note

The enable LIP reset, enable LIP full logon, and enable target reset parameters control the behavior of the adapter when Windows NT tries to do a SCSI bus reset. You must perform a target reset to make cluster failovers work. Use SCSI bus device reset to clear SCSI reservations. The SAN Data Gateway does not support LIP reset and a full login is not necessary after the target reset.

7. Press escape to return you to the Configuration Settings menu.
8. From the Configurations setting menu, scroll down to Extended Firmware Settings menu. Press Enter.
9. From the Extended Firmware Settings menu, scroll down to Connection Options to open a window for Option and Type of Connection.
10. Select the option:

    0 - Loop only

    1 - Point-to-point

    2 - Loop preferred (If you cannot use arbitrated loop, then default to point-to-point)

    3 - Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop).

Note

The option you select must match the port connections on the ESS. Use the StorWatch ESS Specialist to set the Fibre Channel topology on the ESS.

11. Use the down arrow key on the keyboard to highlight LUNs per target.
12. Use the down arrow key on the keyboard to find and highlight 256. Press Enter.
13. Press the escape key.
14. Save the changes. Highlight Yes.
15. Restart the server.

3.1.3 Loading the current Fibre Channel adapter driver

Perform the following steps to load the current driver onto the QLogic adapter card.

1. Go to this Web site: `http://www.qlc.com`
2. From the home page, click Driver Download.
3. Click Drivers.
4. Click Fibre Channel Adapter Drivers.
5. Click QLA2 xxxdrivers.
6. Click Windows NT 4.0.
7. Click Driver.
8. In the Save As pop-up window, get the current driver file, `2 xxxxxxx.exe`.

Note

Save the file to an existing file folder or create a temporary directory where you can save the current driver file.

9. Click Save.
10. Close the Web site.
11. From your Windows NT Start menu, select Run.
12. In the Run pop-up window, ensure the drive letter in the field is the same as the drive letter where you saved the `2 xxxxxxx.exe` file in step 8. If no drive letter appears, type the letter of the drive where you saved the driver file.

13. Type the driver file name after x:\, where x is the drive letter you specified to save the file.
14. Type the directory name where you want to put the file. Click on Zip.
15. Click OK to unzip the current driver file.

3.1.4 Installing the drivers

Perform these steps to install the Fibre Channel adapter drivers:

1. From your Windows NT desktop, double-click the icon for My Computer.
2. Double-click the icon for Control Panel.
3. Double-click the icon for SCSI Adapters.
4. In the SCSI Adapters window, click the tab for Drivers.
5. Click Add.
6. In the Install Drivers window, click Have Disk.
7. In the Install from Disk window, ensure the drive letter in the field for Copy Manufacturer's Files From is the drive letter you specified to save the `2xxxxxxx.exe` file in step 8 on page 56.
8. Type the name of the current driver file in the field for Copy Manufacturer's Files From after the drive letter prompt.
9. Click OK.
10. Click OK to exit.
11. Restart your host system.

3.1.5 Locating the World Wide Port Name (WWPN) on Windows NT

Perform these steps to locate the World Wide Port Name (WWPN) for the Qlogic 2100F adapter card and the Qlogic 2200F adapter card:

1. Restart the server.
2. Press and hold the Alt key and press the Q key to get to the FAST!Util command panel.

   If you have more than one Fibre Channel adapter installed, you will see a window that displays all the Fibre Channel adapters. Scroll down to the adapter you want and press Enter.

   a. From the Fast Util! menu, scroll down and highlight Select Host Adapter.

   b. Scroll up and highlight Configuration Settings. Press Enter.

c. From the Configuration Settings menu, select Host Adapter settings.

d. Write down the host adapter name, for example: `200000E08B00C2D5`.

e. Scroll down and highlight Select Host Adapter settings.

3.2 Installing Fibre Channel host adapter cards on a RS/6000 system

In the following sections, we cover these topics:

3.2.1 The 2105 host attachment package

This section provides the instructions to install the host attachment package for the ESS on your host system. IBM recommends that you run the host attachment package on each host system attached to the ESS for which an installation script is provided.

Note

For an RS/6000 host system, you can use either point-to-point (switched fabric) topology or an arbitrated loop (hub) topology.

3.2.2 Before you install the 2105 host attachment package

Before you install the host attachment package:

1. Attach the ESS to your host system.
2. Turn on the host system and all attachments.

Note

Before you install the host attachment package, you must have:

- Root access
- AIX system administrator knowledge
- Knowledge of the Software and Management Interface Tool (SMIT) to install the ESS host attachment package

3.2.3 Replacing an older version of the 2105 installation package

If you want to replace the older version of the host attachment package (tar version) and have data that exists on all configured 2105 disks, the code prompts you to remove all ESS product-related hdisk devices. Perform the following steps to remove the devices:

1. Run umount on the file system.
2. Run varyoffvg 2105 volume group.
3. Type `rmdev -dl` on the command line to unconfigure the 2105 devices.

After you install the ibm2105.rte file and all of the 2105 devices are reconfigured, vary on the volume groups and remount the file systems. The data on the file systems should now be available again.

Perform the following steps by using SMIT to install the IBM 2105 host attachment on your system.

3.2.4 Installing the 2105 host attachment package

Note

The following procedure is an example. The example uses /dev/cd0 for the address of the compact disc. Your address may vary.

Install the host attachment package from a compact disc or a diskette. You must have superuser authority to complete the instructions.

1. From your desktop window, type `smit install_update` to go directly to the installation panel.
2. Click Install and Update from the Latest Available Software and press Enter.
3. Press F4 to display the Input Device / Directory for Software window.

4. Select the CD-ROM drive that you are using for the installation, for example, /dev/cd0. Press Enter. The Install and Update from the Latest Available Software window displays.
5. Click Software to Install and press F4. The Install and Update from the Latest Available Software panel displays with the name of the software you selected to install.
6. Check the default option settings to ensure that they are what you need.
7. Press Enter to install the software. SMIT responds with the following question: `Are you sure?`
8. Press Enter to continue. The installation process may take several minutes. A message displays on your window when the installation process is complete.
9. Verify whether or not the installation is successful.
10. Press F10 when the installation process is complete.
11. Exit from SMIT.
12. Remove the compact disc.
13. Shut down the host system.
14. Turn on the host system.

3.2.5 Locating the World Wide Port Name (WWPN) on RS/6000

Perform this step to locate the WWPN:

1. Log in as root. Type `lscfg -vl fcsx`, where x is the adapter number. The network address is the Fibre Channel adapter port WWPN value.

3.2.6 Verifying the IBM ESS configuration

To verify the configuration of the ESS on the AIX host system, type the following command:

```
lsdev -Cc disk | grep 2105
```

The installation is successful if a list of all IBM ESS devices displays:

```
hdisk3 Available 30-68-01 IBM FC2105F20
hdisk4 Available 30-68-01 IBM FC2105F20
hdisk5 Available 30-68-01 IBM FC2105F20
...
...
```

If a device is listed as another type of device, this message displays. This message indicates that the configuration was not successful:

```
hdisk3 Available 30-68-01, Other FCSCSI disk device
hdisk4 Available 30-68-01, Other FCSCSI disk device
hdisk5 Available 30-68-01, Other FCSCSI disk device
...
...
```

When you use the `lsdev -Cc disk | grep 2105` command, you see only display lines that contain the value immediately after it. If you have not defined any 2105 devices, none will be displayed.

Chapter 4. Configuring the ESS with native Fibre Channel

One benefit of a SAN is to implement disk pooling. To do this successfully, we need an easy to manage storage server with a Fibre Channel attachment. This is the IBM Enterprise Storage Server, 2105-F20 with native Fibre Channel.

The prime focus of this redbook is directed towards the fabric, rather than the devices connected to it. We describe what we need to do to make the ESS part of our SAN.

The following hardware and software, necessary for successful attachment, should be installed before beginning to set up the SAN:

- Fibre Channel adapters installed in the ESS
- Microcode level on the ESS (1.1) to support the Fibre Channel attachment
- InfoServer must be available

In the topics that follow, we cover:

4.1 Defining the ESS

As with the SCSI-ESCON ESS, we will use the StorWatch Enterprise Storage Server Specialist to configure the hosts, ports, and volumes to be accessed from the hosts. We only describe what we have to configure for use with FC-AL and FC-SW environments. For further information on the configuration of the Fibre Channel ESS refer to the ESS documentation, as described in 4.2, “Related information” on page 109.

In the following sections, we cover these topics:

4.1.1 Introduction to the ESS Specialist

The ESS Specialist is a Web and Java based interface to configure the ESS. To connect to it, we use the Netscape Navigator and use the hostname or IP address of the ESS in the Uniform Resource Locator (URL) field.

4.1.1.1 Introduction panel

This takes us to the introduction panel of the ESS Specialist, as shown in Figure 25.

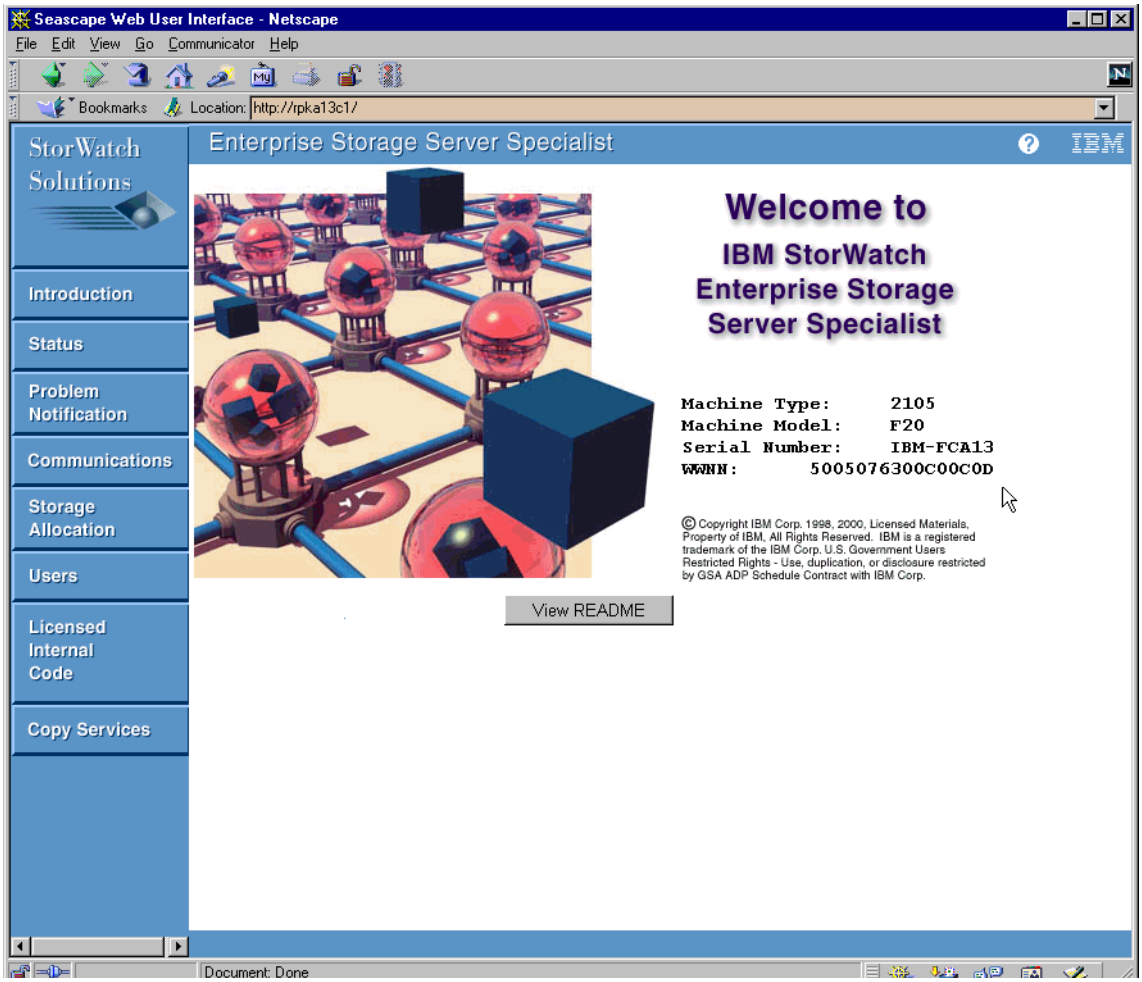

Figure 25. Introduction panel of the native FC ESS

Here we can see the machine type, model, and serial number displayed. Also we see the identification of the Fibre Channel ESS in the Fibre Channel fabric in the form of its WWNN.

To the left, we see the main button panel, from where we start to configure the options for the ESS. The Introduction button will load the Introduction panel:

4.1.1.2 Logon to the ESS

Using one of the buttons to the left, we are forced to login to the site with the login window shown in Figure 26. The default user is `storwatch` and the password is `specialist`. Both are case sensitive and must be entered in lower case.

Figure 26. Login window

4.1.1.3 Status Panel

Using the Status button we get to the Status -- Graphical View panel, as shown in Figure 27.

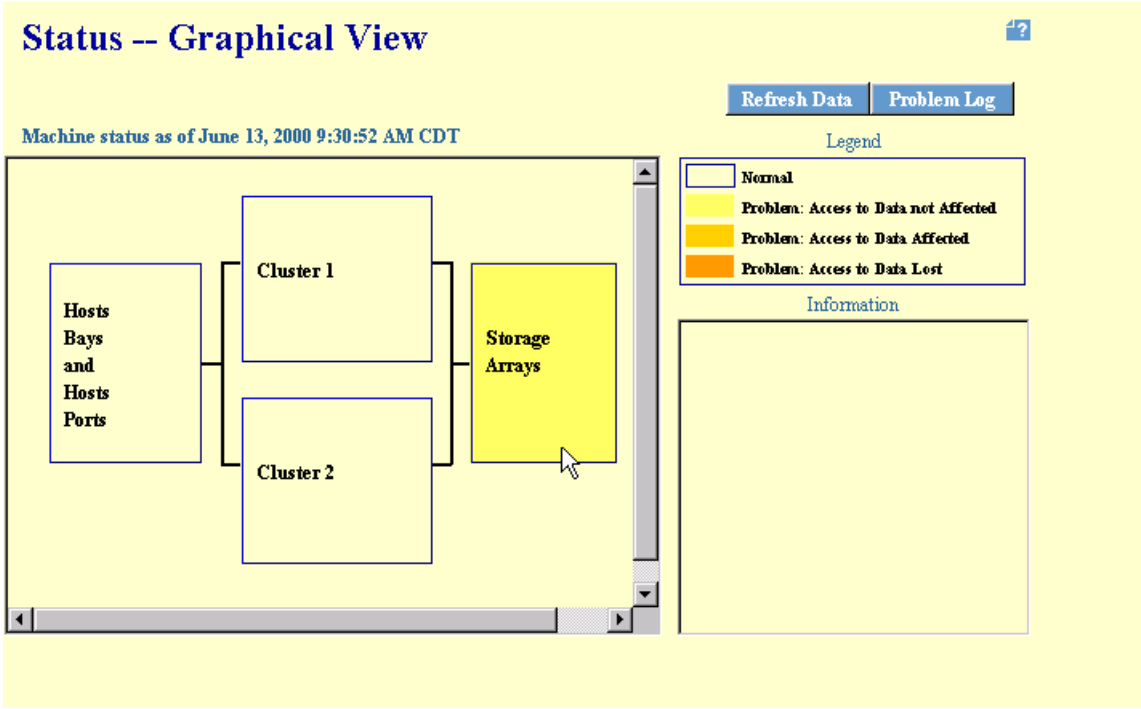

Figure 27. Status panel graphical view

This panel allows us to quickly check for any errors in the ESS. Here, we see a problem detected in the storage arrays, with access to data not affected. To get a detailed view on the problems encountered, we use the Problem Log button and get the Status -- Problem Log panel, as shown in Figure 28.

Figure 28. Status panel problem log

4.1.1.4 The Problem Notification panel

Using the Problem Notification button we get the Problem Notification panel, shown in Figure 29.

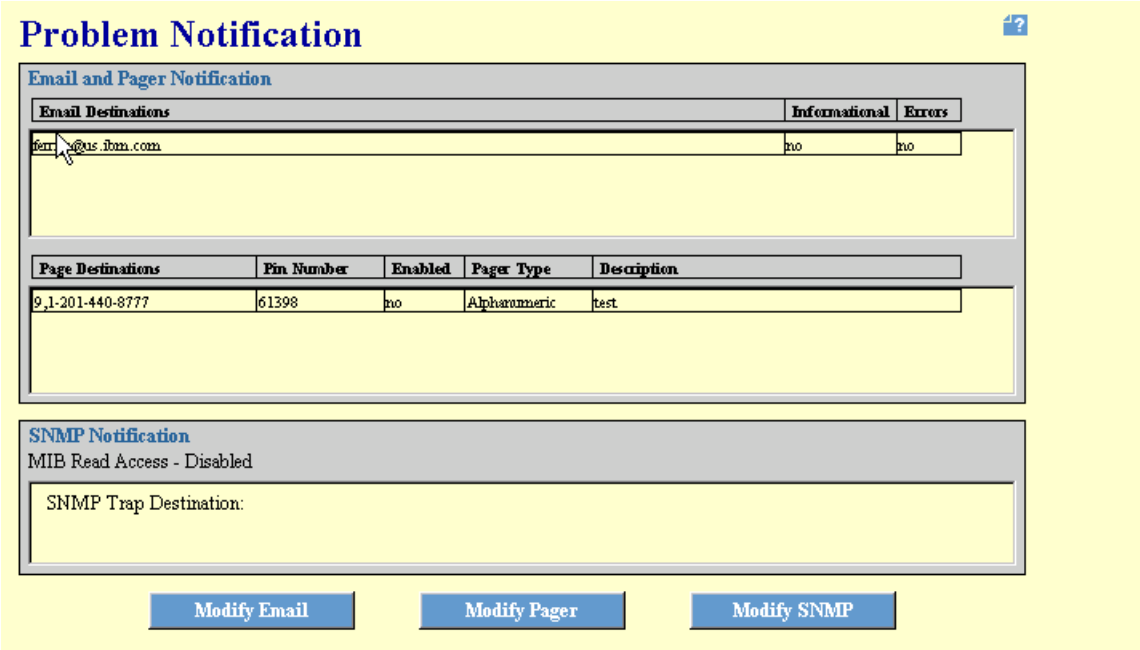

Figure 29. Problem notification panel

From within this panel we can configure the notification options provided with the ESS. Also, we can configure the Simple Network Management Protocol (SNMP).

4.1.1.5 The Communications panel

Using the Communications button we get the Communications panel, as shown in Figure 30.

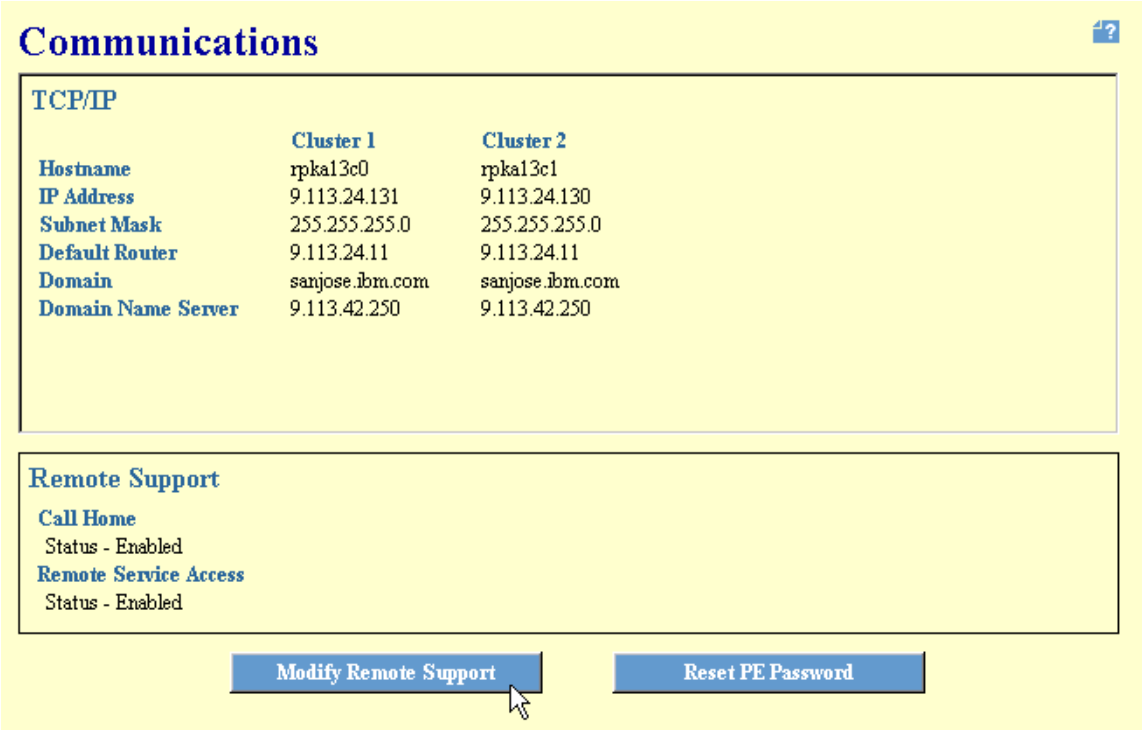

Figure 30. Communications panel

Here we can view the network information, for which the ESS is configured, and verify if the call home feature and remote service access is enabled. Also we can change the remote support features, shown in Figure 31, and reset the PE password.

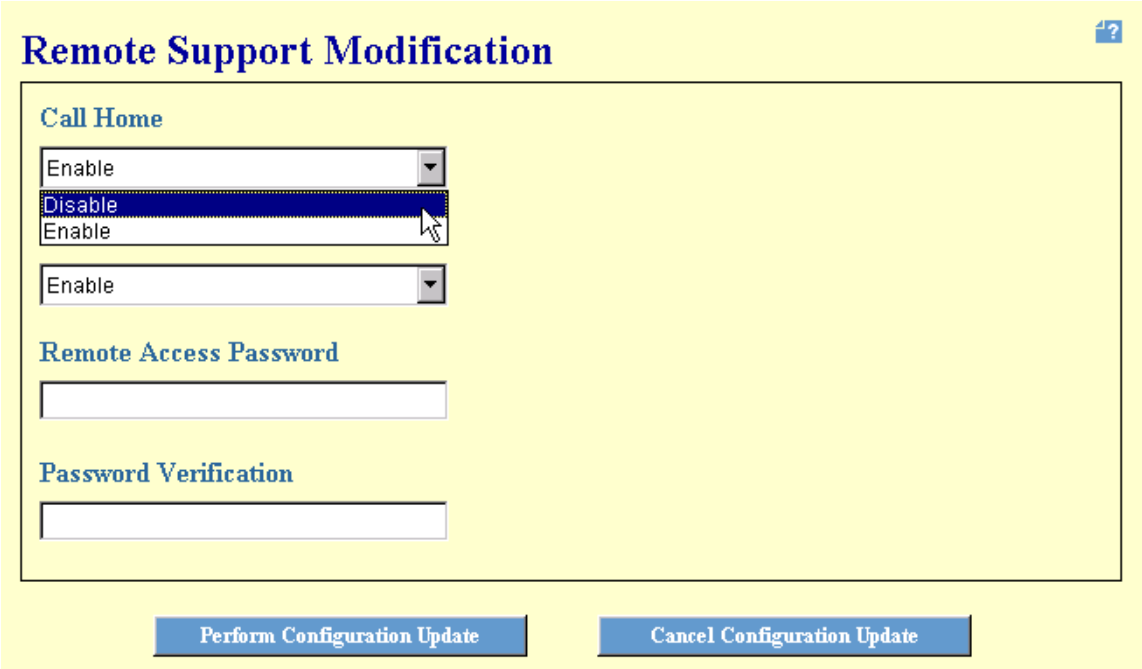

Figure 31. Remote support modification

4.1.1.6 Storage allocation

Using the Storage Allocation button we get the Storage Allocation -- Graphical View panel, shown in Figure 32.

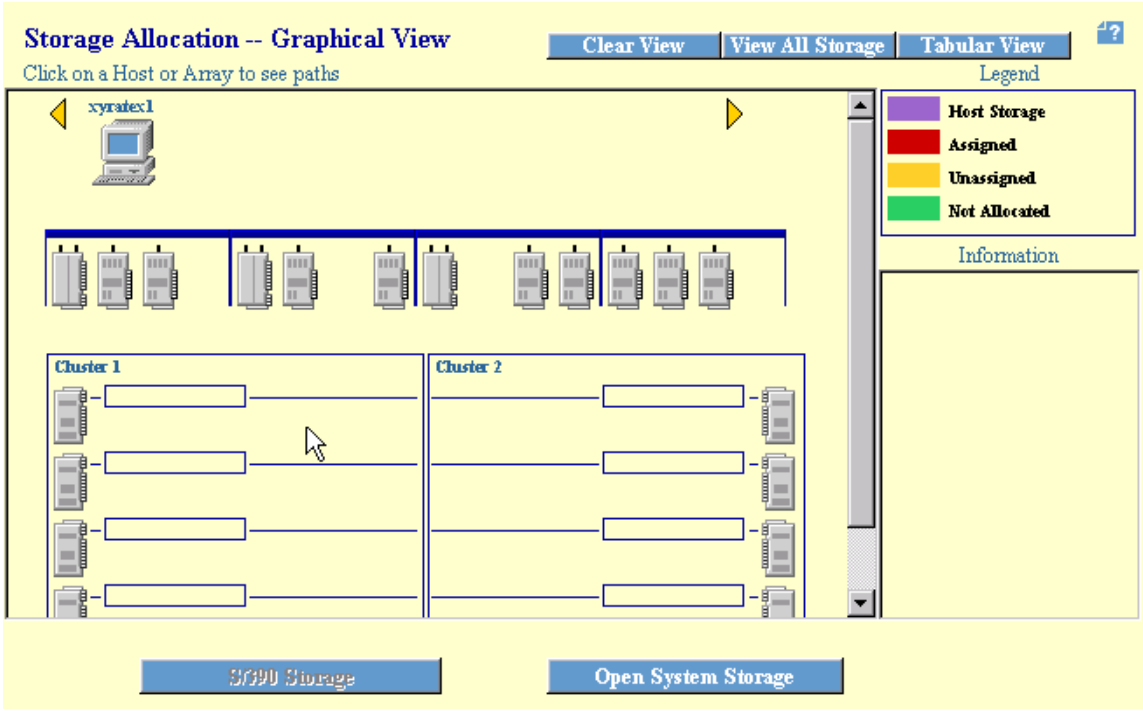

Figure 32. Storage Allocation panel

This is the start panel for host and storage based configurations. In 4.1.2, “Viewing the Storage Allocation of the ESS” on page 75, we will show how to configure the ESS for our SAN.

For completeness we will show the other panels.

4.1.1.7 User Administration

Using the Users button we get the User Administration panel, as shown in Figure 33.

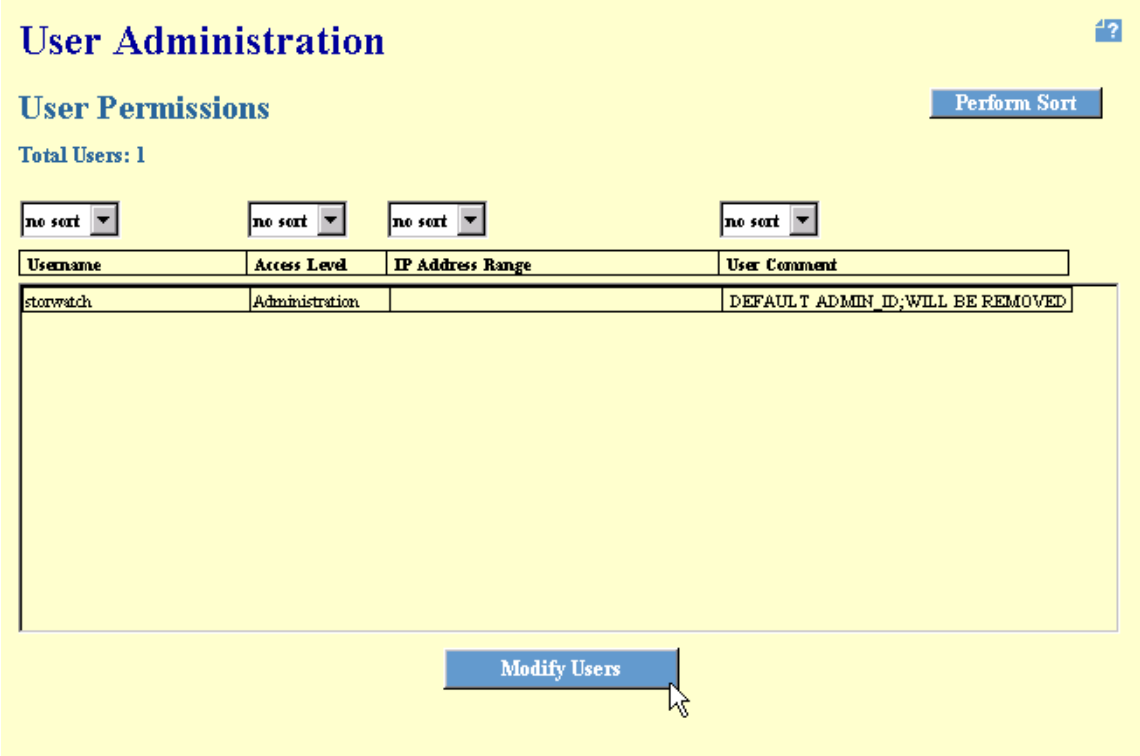

Figure 33. User administration panel

In our case there is no user configured, other than the default user that we logged on with. In the Modify Users panel, as shown in Figure 34, we can add users with different access levels.

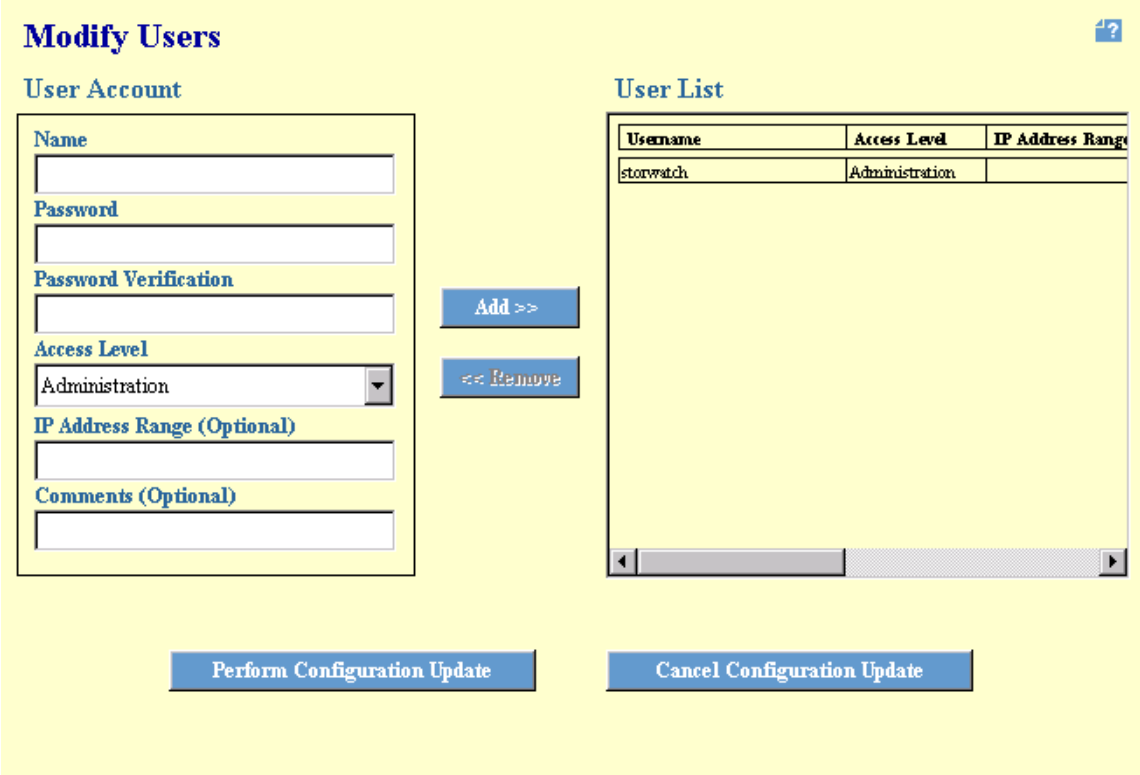

Figure 34. Modify users panel

4.1.1.8 Licensed Internal Code

Using the Licensed Internal Code button we get the ‘Licensed Internal Code’ panel, as shown in Figure 35.

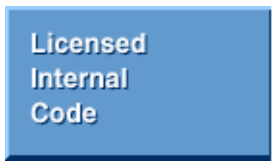

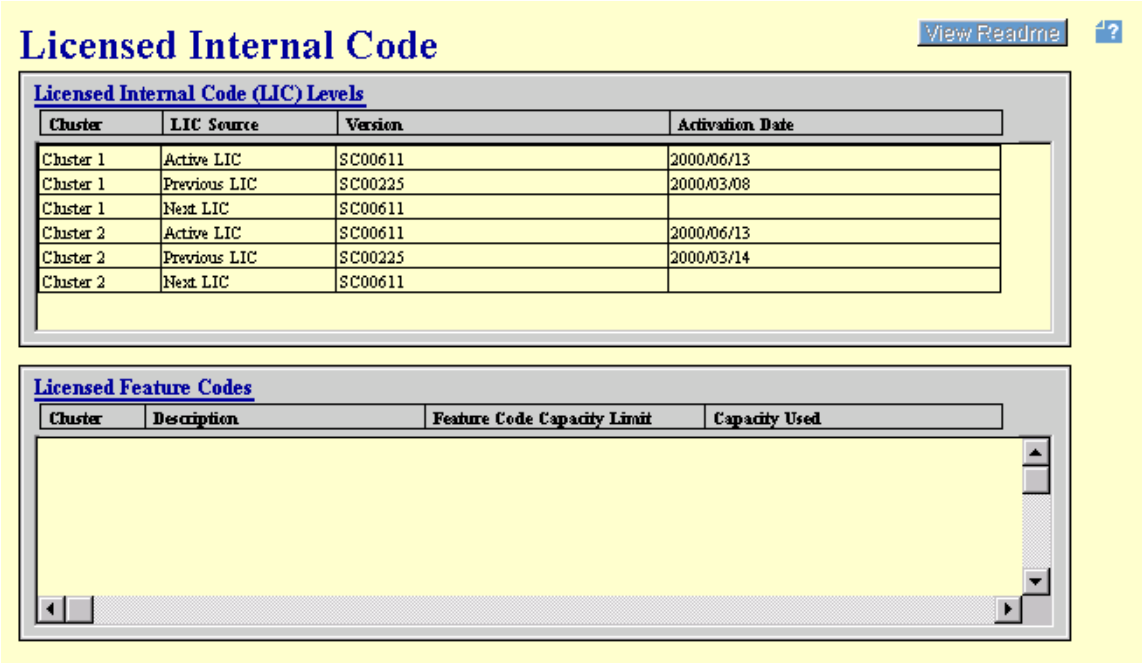

Figure 35. Licensed Internal Code panel

Here we can see the LIC levels and the licensed feature codes.

4.1.1.9 Copy Services

Using the Copy Services button we connect to the copy server.

We have to login to gain access, just like we did for the ESS Specialist. This opens the Web based Java interface for the ESS Copy Server in a dedicated Navigator window, and connects to the copy services server, as shown in Figure 36.

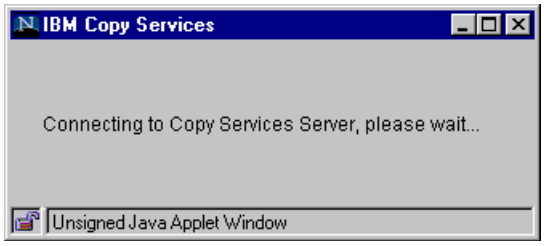

Figure 36. Connecting to Copy Services Server

Once connected, we get the introduction panel for all copy service configurations, as shown in Figure 37.

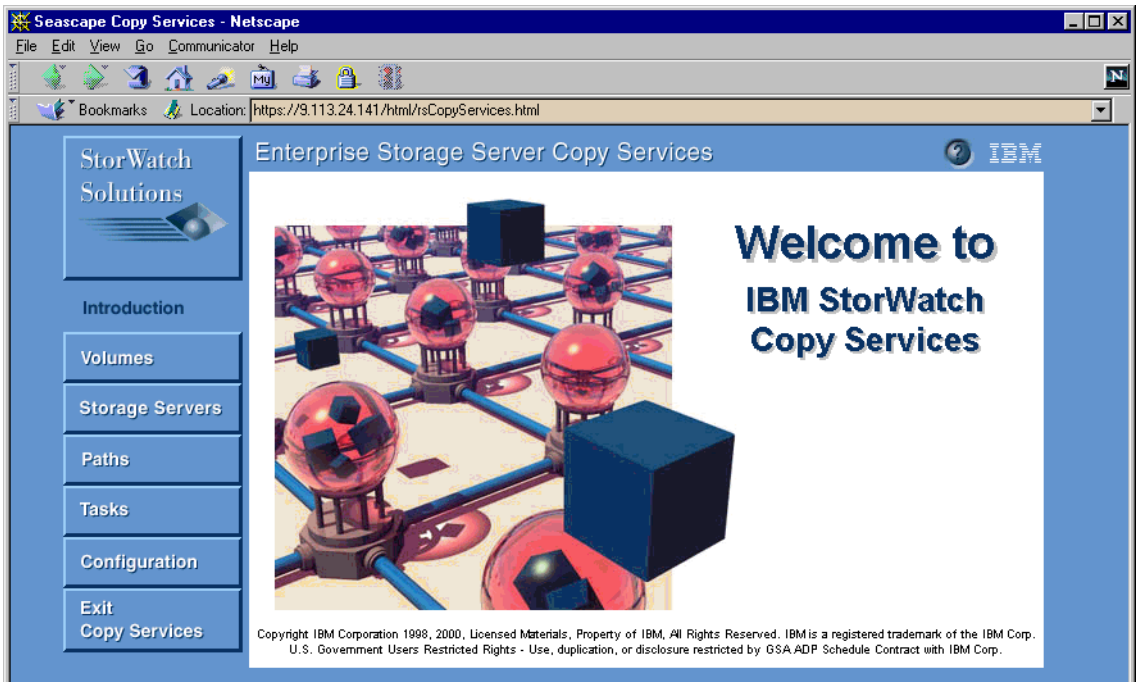

Figure 37. ESS Copy Services introduction panel

Copy Services will be covered in future IBM Redbooks.

4.1.2 Viewing the Storage Allocation of the ESS

Using the Storage Allocation button we get the Storage Allocation -- Graphical View panel, as shown in Figure 38.

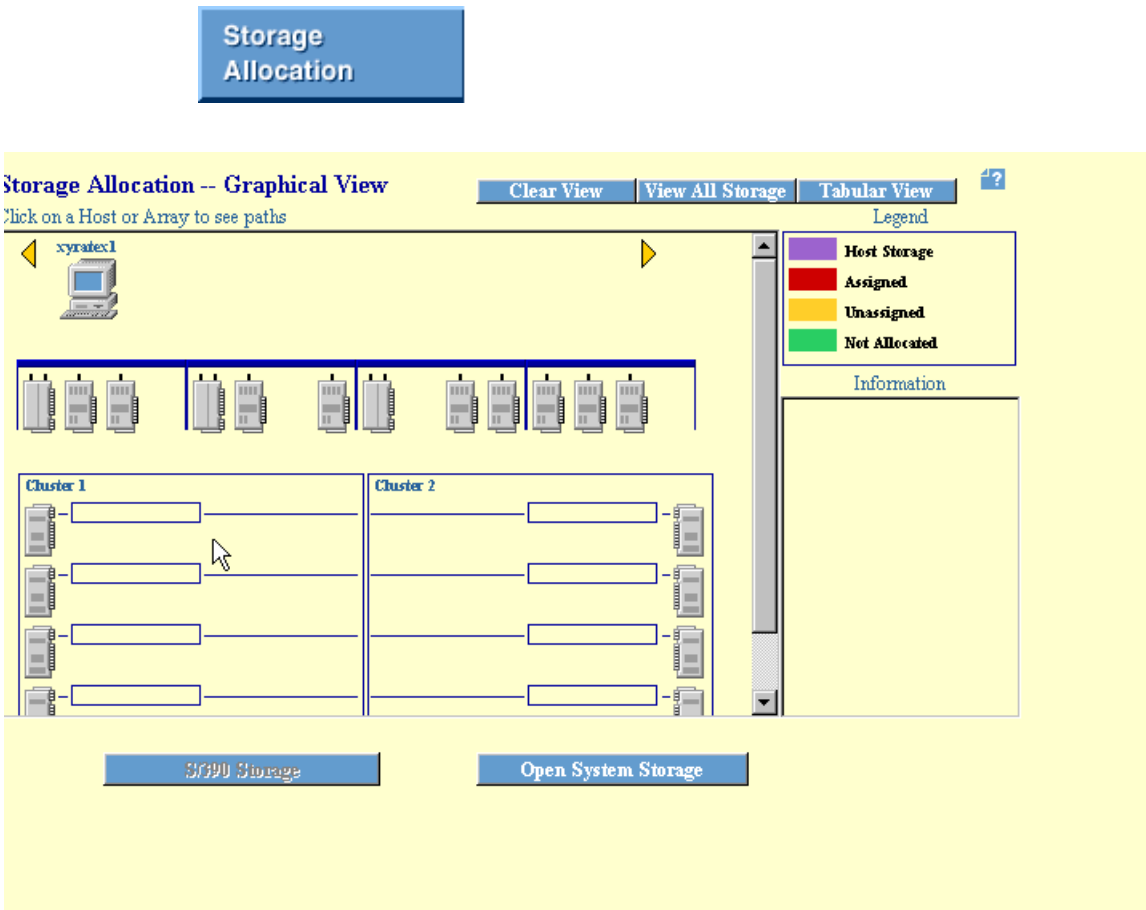

Figure 38. Storage Allocation ESS with FC and SCSI

We can see that our ESS is equipped with nine Fibre Channel adapters and three SCSI adapters. Graphically, the representation of SCSI adapters (with two ports per adapter) is shown in Figure 39.

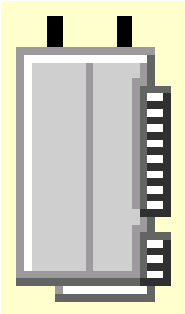

Figure 39. SCSI adapter with two ports

The Fibre Channel adapter with one port is graphically represented as shown in Figure 40.

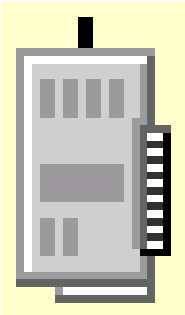

Figure 40. FC adapter with one port

In the ESS, we do not tie a host to a Fibre Channel port in the ESS, which was what we did with the SCSI adapters. Every host adapter port within the Fibre Channel fabric will be identified by its World Wide Port Name (WWPN) and can access data through every ESS Fibre Channel WWPN unless other measures are taken.

To show the independence of the port, the selected host is graphically connected to each Fibre Channel adapter, as shown in Figure 41.

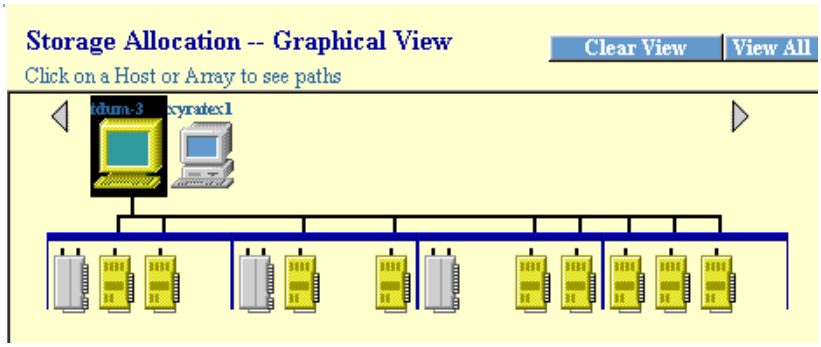

Figure 41. FC host selected

In contrast to this, the SCSI attached hosts have an affinity to the SCSI ports on the ESS, which is visible by selecting a SCSI host, as shown in Figure 42.

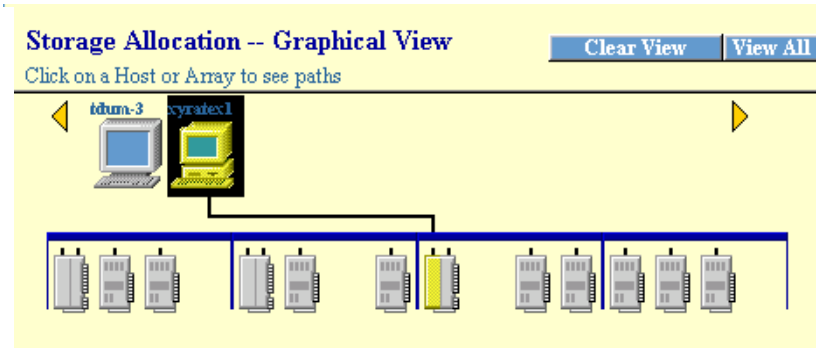

Figure 42. SCSI host selected

To get a detailed view of the assigned volumes, we use the Tabular View button to get to the view as shown in Figure 43.

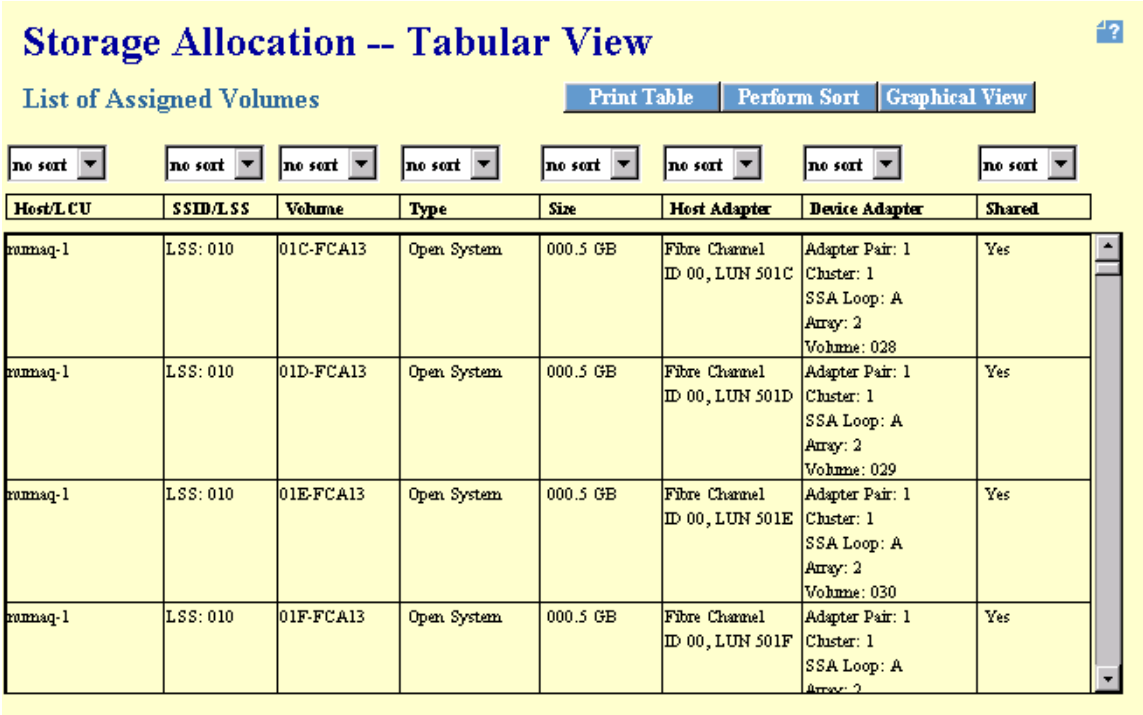

Figure 43. Storage Allocation tabular view

To perform the configuration tasks that we need, we press the Open System Storage button, as shown in Figure 44.

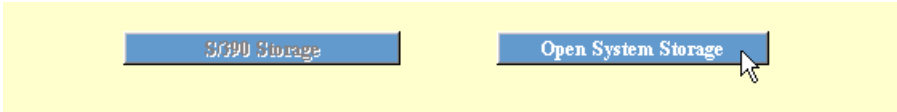

Figure 44. Entry for open system storage

The S/390 Storage button is grayed out, because there are no ESCON cards installed.

4.1.3 Accessing the Open System Storage panel

This will lead us to the Open System Storage panel, as shown in Figure 45. The buttons at the bottom of the panel are the entry points to all of our host and storage configuration tasks. In this section, we will indicate our use of those buttons by pointing the mouse on the button.

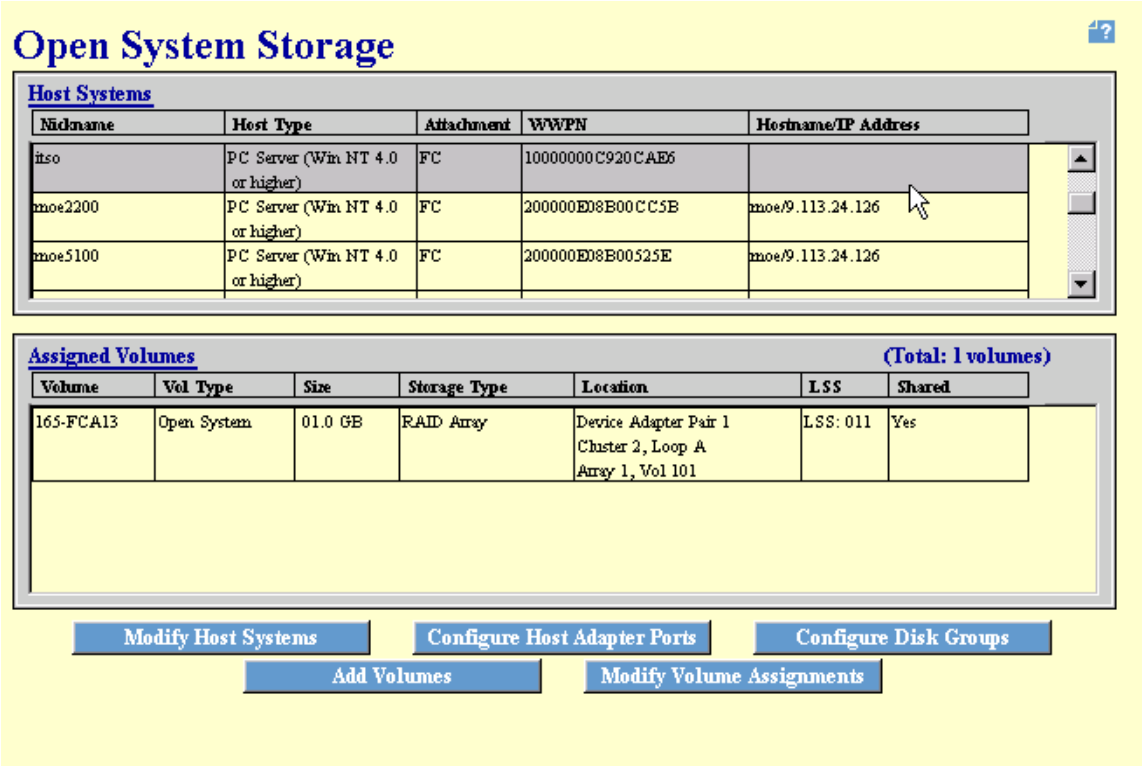

Figure 45. Open System Storage panel

Here we can see the Hosts Systems table and the Assigned Volumes table. We will see shortly that those two tables are showing the defined host FC adapter ports and, to a selected port, the assigned volumes.

4.1.4 Defining a new host FC port with its WWPN

We pressed the Modify Host Systems button, as shown in Figure 46.

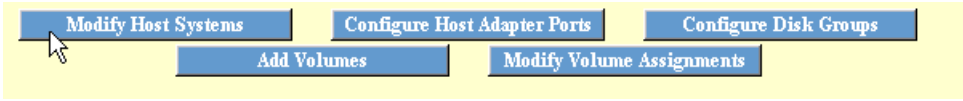

Figure 46. Entry to the Modify Hosts Systems panel

We are taken to the Modify Host Systems panel, as shown in Figure 47. This panel consists of the Host Attributes entry fields and The Host Systems List table. The Host Systems List represents each hosts' Fibre Channel port. If a host has two Fibre Channel adapter cards, this appears as two hosts.

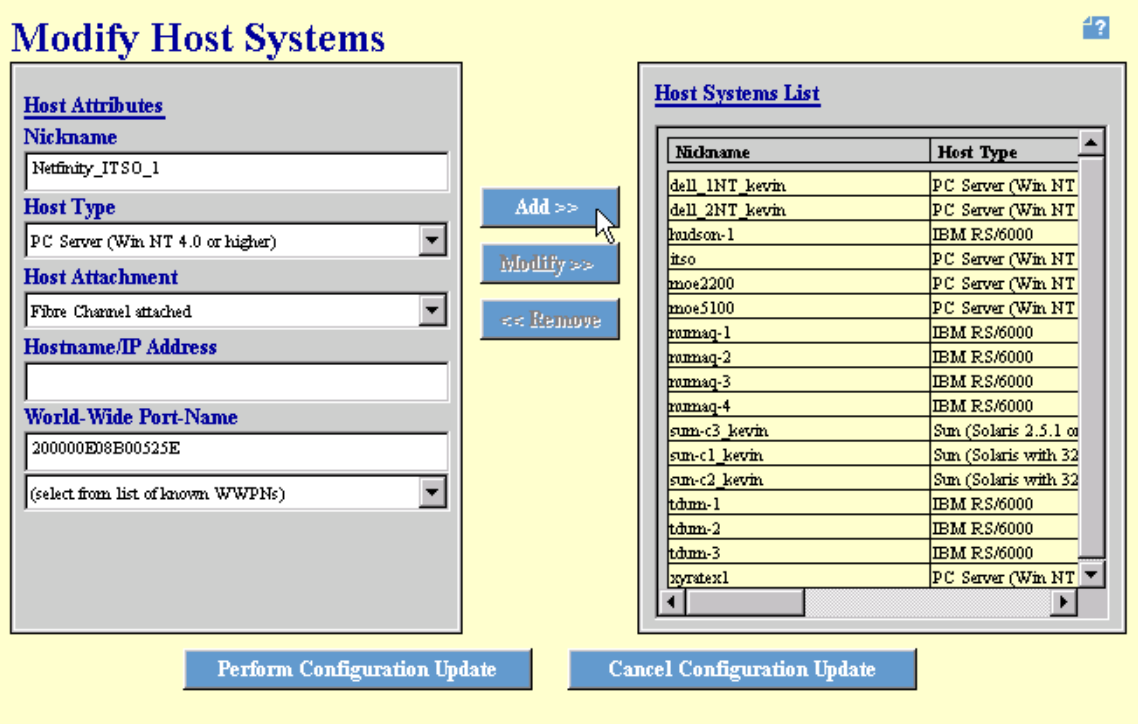

Figure 47. Modify Host Systems panel

In the Host Attributes entry fields, we identify the host Fibre Channel ports that will access the ESS. The first entry is a nickname to identify the port in the ESS. Also, we have to choose the host type, which is a PC server; the type of the host attachment, which is Fibre Channel; and the WWPN of the host adapter port to identify the host adapter. If you plan to use the StorWatch Enterprise Storage Server Expert later on, we recommend that you also enter the host IP address field. This is used to show which Fibre Channel host adapters are in a specific IP host.

By pressing the Add button, we add the specified host Fibre Channel port (of a Fibre Channel host) to the Host Systems List, as shown in Figure 48.

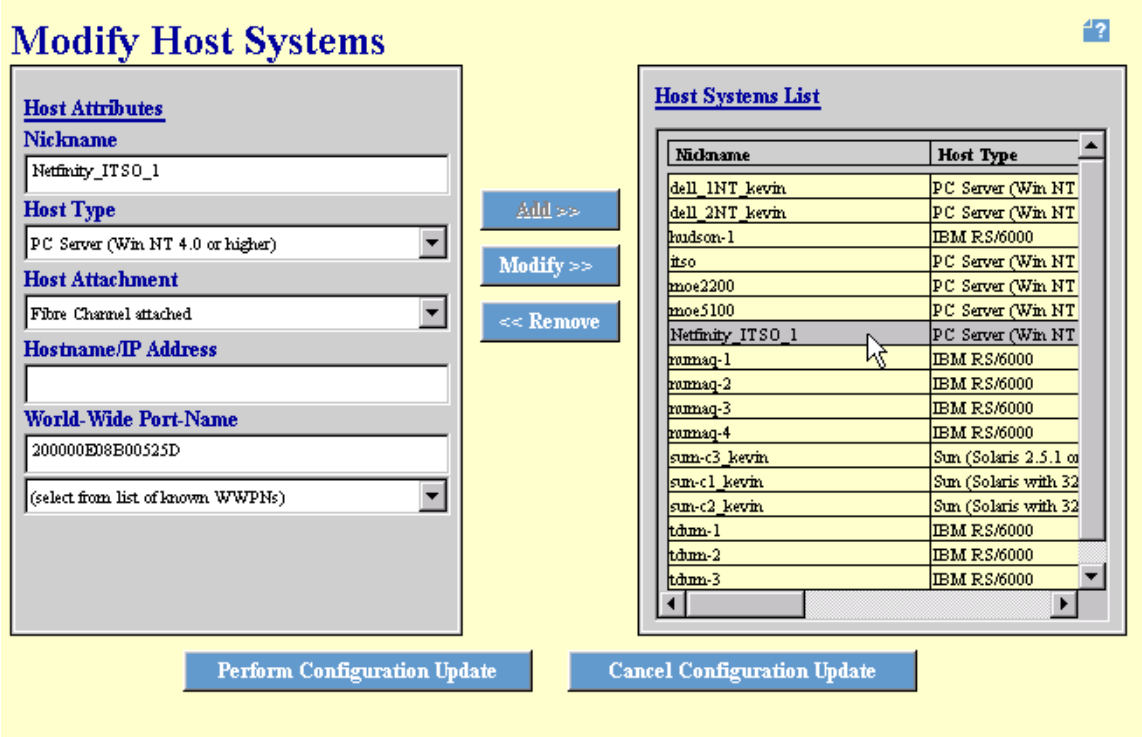

Figure 48. Adding a FC host adapter port

If we wanted to, we can define a second host FC port belonging to the same system, or we can define ports for other systems. Adding ports to the list puts together a script in the background which is executed by pressing the Perform Configuration Update button.

A progress bar informs us about the steps that are being executed, as shown in Figure 49.

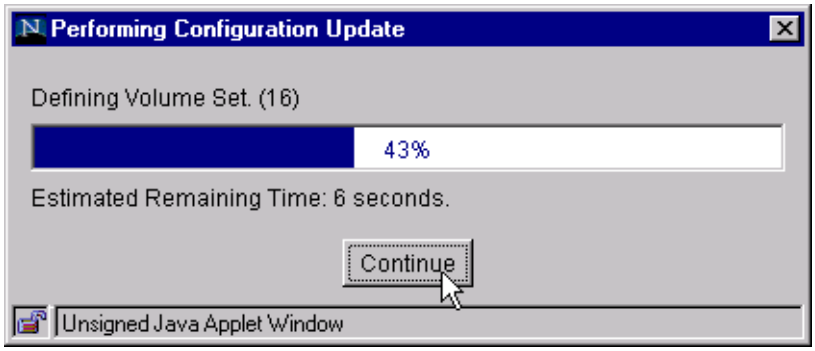

Figure 49. Performing the port definition

A window appears to indicate that the configuration update was successful, as shown in Figure 50. If the update had failed for any reason, we will get an error message.

Figure 50. Successfully added host FC port

Pressing the OK button on this window will immediately take us back to the Open System Storage panel, where we now see the newly configured host Fibre Channel port in the Host Systems Table, as shown in Figure 51.

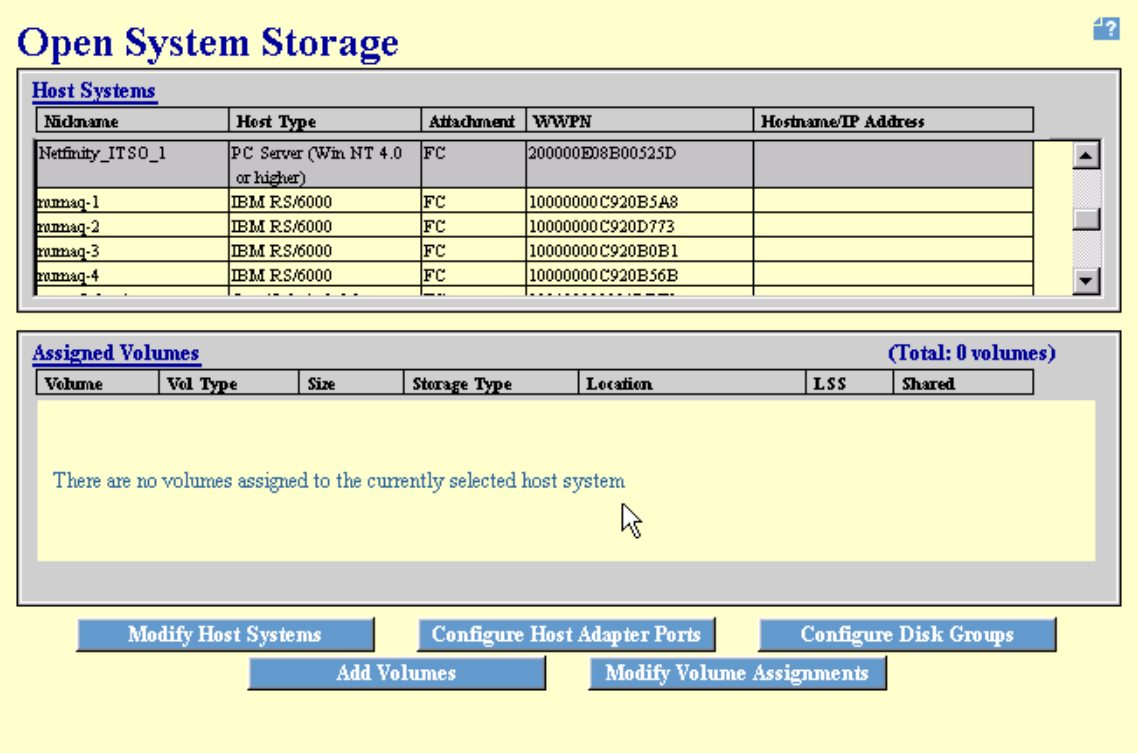

Figure 51. Newly configured Fibre Channel host adapter port; no volume assigned

At this point, we can work with other configurations, for instance, configuring the options for the ESS Fibre Channel ports we want to use, or assigning volumes to the newly defined host Fibre Channel port.

To configure disk groups and volumes from the beginning, even in a partially configured ESS, we first want to configure disk groups.

4.1.5 Configuring disk groups

Pressing the Configure Disk Groups button takes us to the Fixed Block Storage panel, as shown in Figure 52.

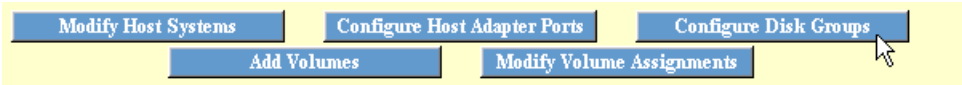

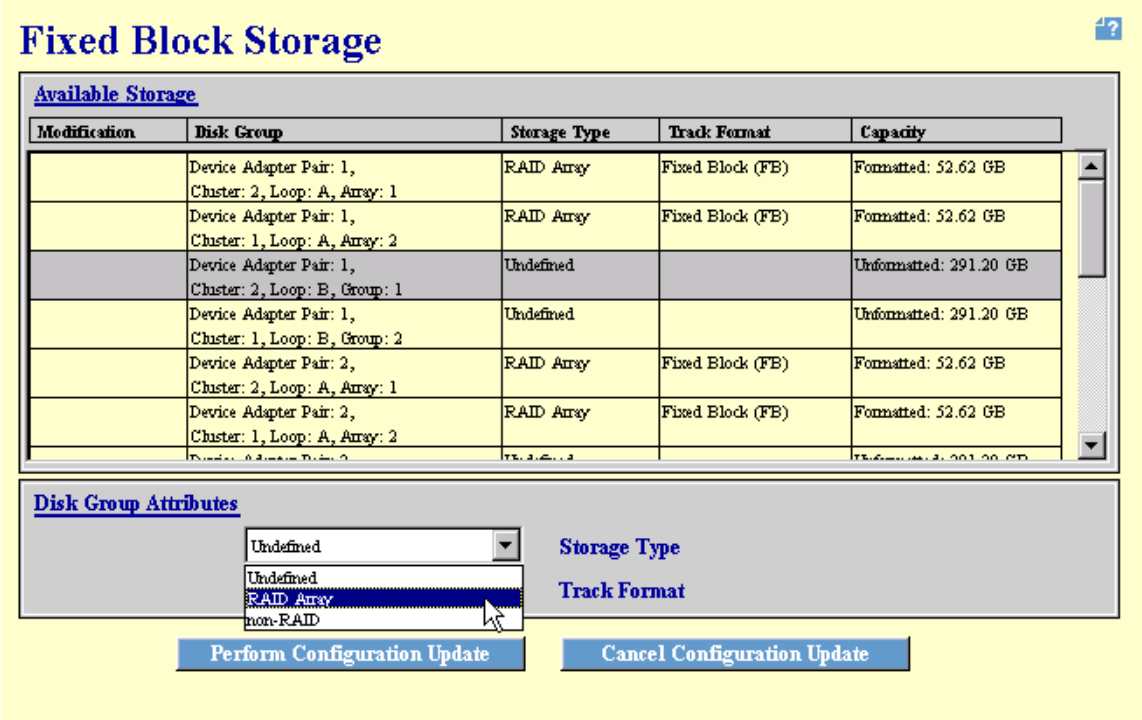

Figure 52. Fixed Block Storage panel

The Available Storage table lists all of the available physical storage in the ESS with their location, storage type, track format, and capacity. Also, modifications made prior to performing an update are indicated.

The storage type can be RAID Array, or non-RAID. Non-RAID means the disks are going to be configured as “just a bunch of disks” (JBOD).

We are going to configure the disk groups connected to the B loop of the first adapter pair as RAID array.

To do this we select the first disk group we want to configure and select RAID array from the Storage Type pull-down menu, and this is also shown in Figure 52. Doing this also changes the Track Format to Fixed Block (FB) in the Track Format pull-down menu. After doing this, we select the second disk group in the other cluster, and also select RAID-array. The Modification column of the panel, as shown in Figure 53, shows that we have made changes (Defined), and the Capacity column shows that our disk groups are still unformatted. In this case, formatted to the ESS means it is not formatted as either a RAID array nor as non-RAID.

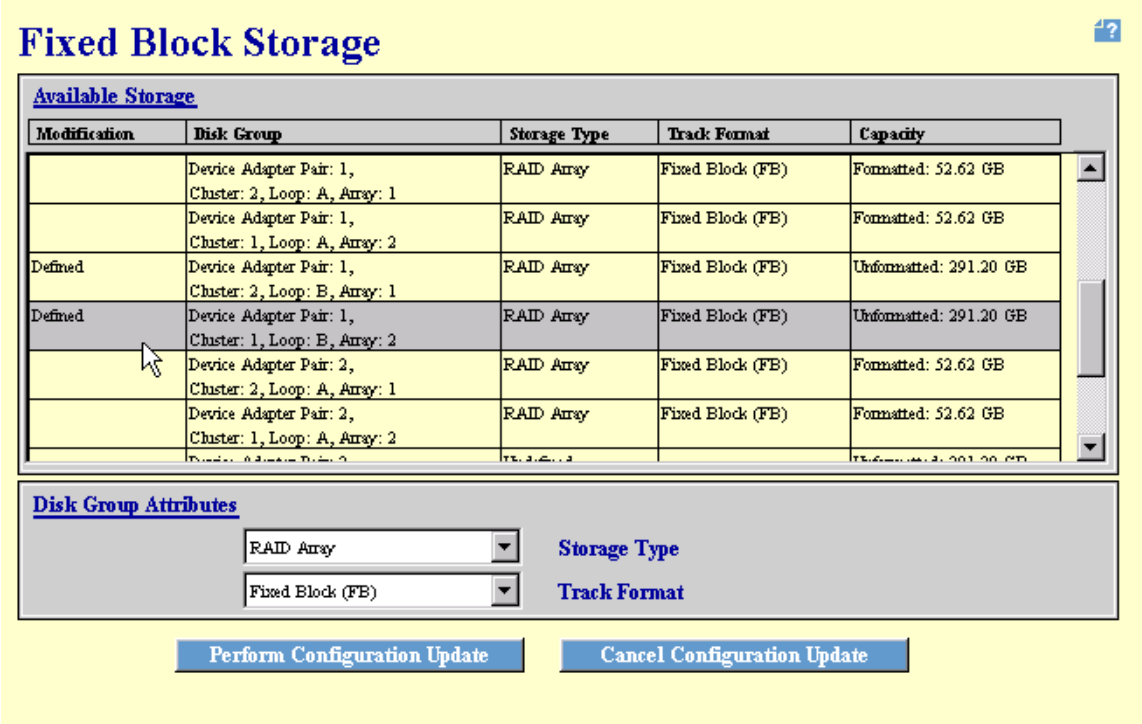

Figure 53. Changed disk group definition

To apply the changes made, we pressed the Perform Configuration Update button.

A warning message appears, stating that this will be a time consuming action, as shown in Figure 54.

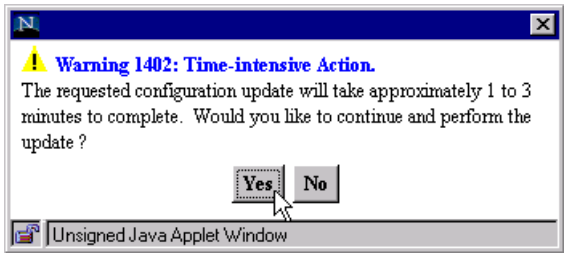

Figure 54. Time consuming action warning

How time consuming this will be depends on how many disk groups are to be configured at once. Pressing the OK button executes the script which changes the configuration and gives us a progress window, where we can see which arrays are initializing, as shown in Figure 55.

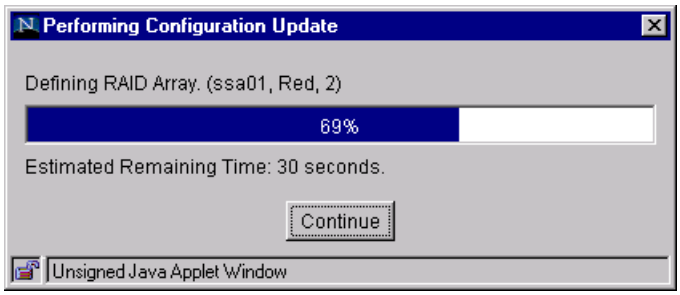

Figure 55. RAID definition progress window

After completing the configuration, the Specialist informs us of the success, as shown in Figure 56.

Figure 56. Successful disk group configuration

We have to press the OK button here, which takes us back to the Open System Storage panel. From there we can continue with our configuration.

If you go back to the Storage Allocation window, device adapter pair number one will now look like Figure 57.

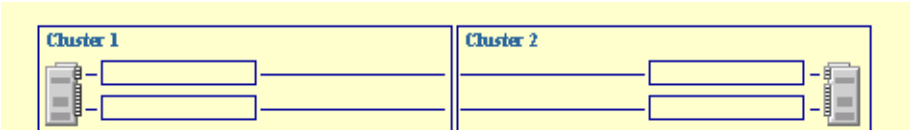

Figure 57. Device adapter pair one with four disk groups

4.1.6 Defining volumes for Fibre Channel host adapter ports

If we return to the Open System Storage panel, we use the Add Volumes button to define volumes on the newly configured disk groups, as shown in Figure 58.

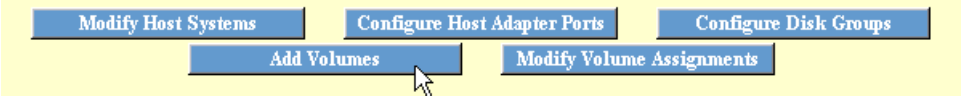

Figure 58. Entry to the Add Volumes (1 of 2) panel

This takes us to the Add Volumes (1 of 2) panel, as shown in Figure 59.

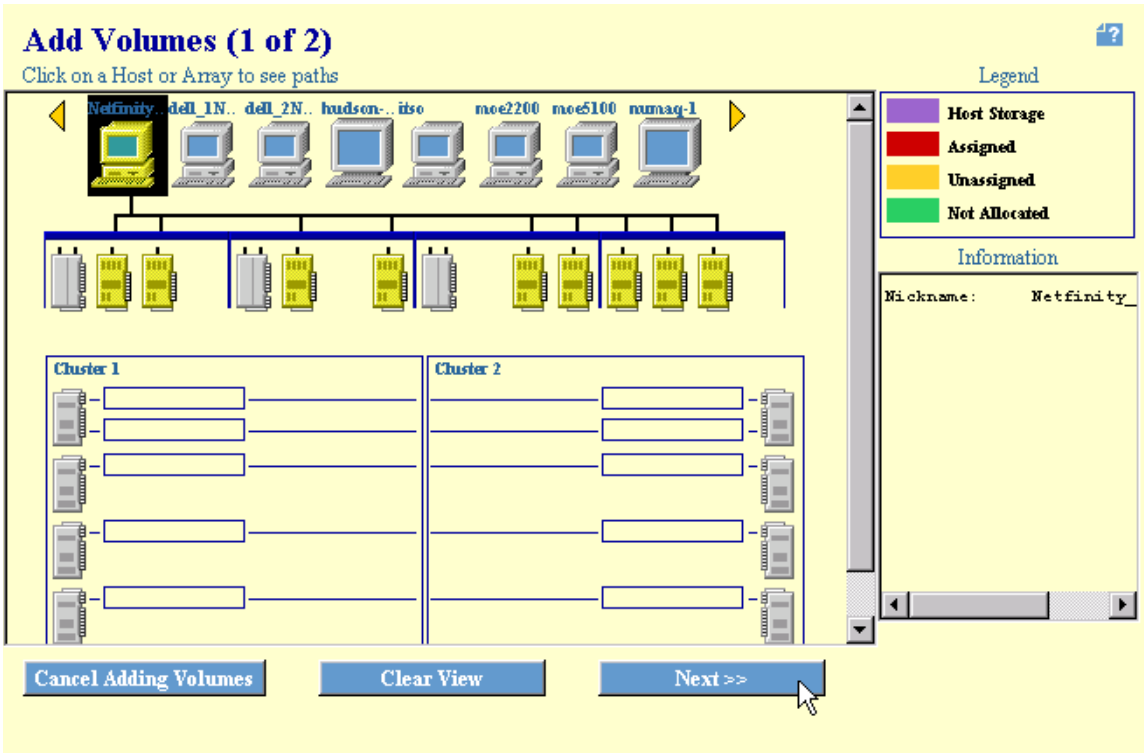

Figure 59. Add Volumes (1 of 2) panel with one Fibre Channel host port selected

From here, we select the Fibre Channel port that we configured. Again, this is *not* the host, but is graphically shown as the host, because our host has only one adapter built in.

We also have to select one of the ESS Fibre Channel ports (Figure 60) before we can go to the next panel.

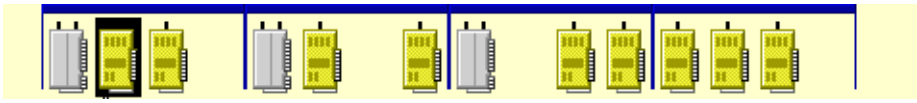

Figure 60. Selecting an ESS Fibre Channel port

Logically, there is no reason to have to select a specific Fibre Channel port, because we can reach the volumes through every properly configured ESS Fibre Channel port. However, by clicking on an ESS Fibre Channel port, we can now see the storage that we can use to define volumes on. This is shown in Figure 61.

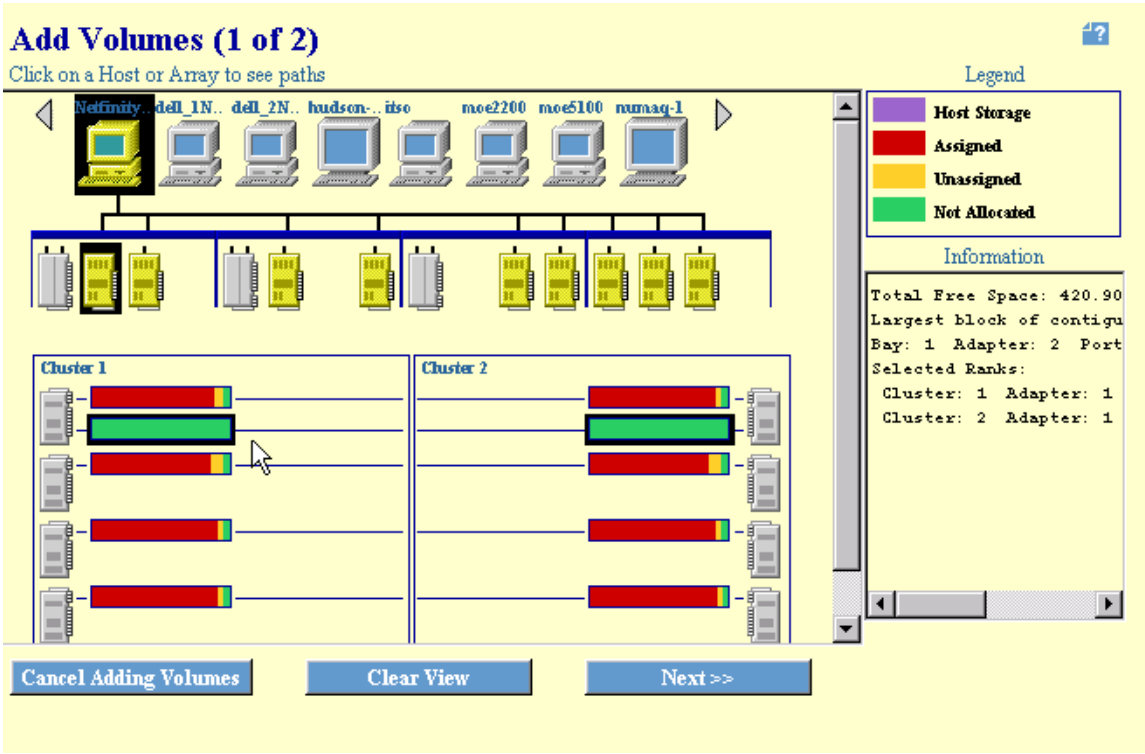

Figure 61. Add Volumes (1 of 2) panel with Not Allocated space selected

There are different colors to indicate which state the disk groups are in. In our case, most of the storage is already configured, only the two disk groups that we configured before have not been allocated. This is indicated with a green color.

To mark the storage for use in the next step, we click on these, which highlights these with a border, as shown in Figure 61.

Then if we press the Next button, we are taken to the second step of defining the volumes.

This is the Add Volumes (2 of 2) panel, as shown in Figure 62.

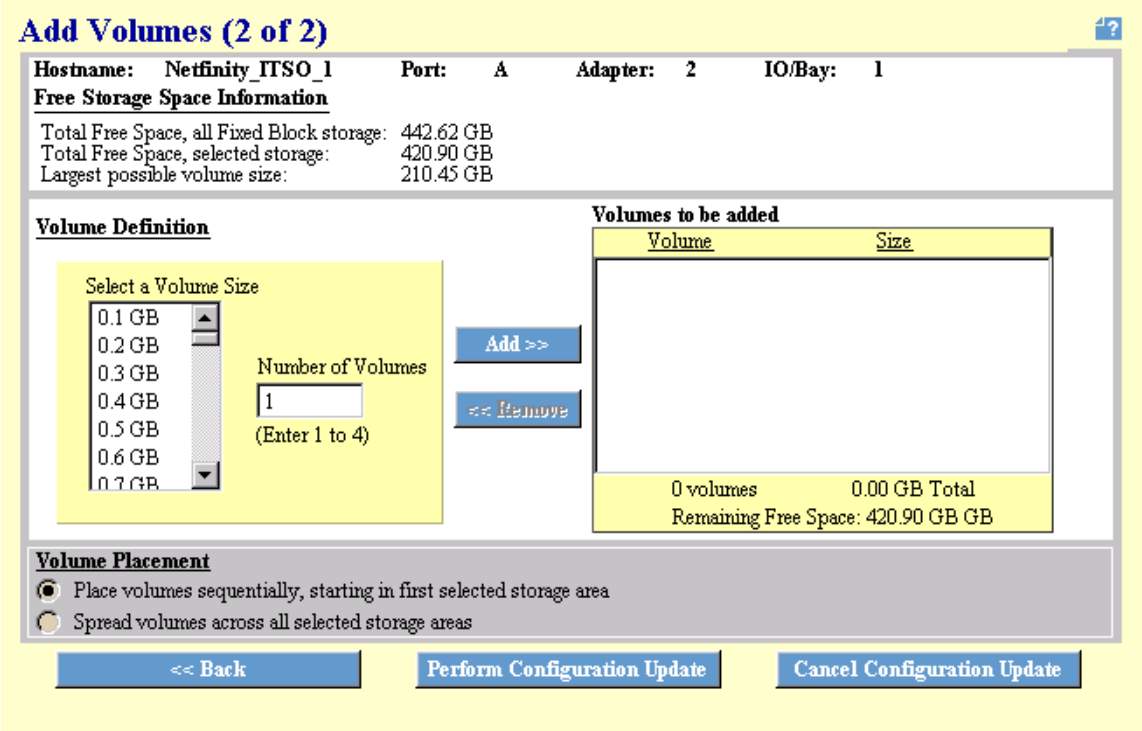

Figure 62. Add Volumes (2 of 2) panel

At the top of the panel, we see the selected Fibre Channel host and port. We can see the total free storage, the size of the storage we selected, and the largest possible volume size.

In the left hand panel, we can choose the size of the volumes that we want to configure. We choose to define four volumes with 96 GB each. This is shown in Figure 63.

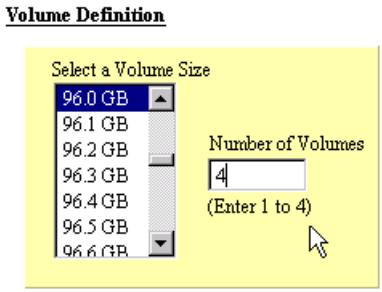

Figure 63. Volume definition

Using the Add button, we place the selected amount of volumes in the Volumes to be added panel.

If we wanted to, we can add other volumes.

We can also choose if we want to place the volumes sequentially, or if we want to spread the volumes over all selected storage areas. This is shown in Figure 64.

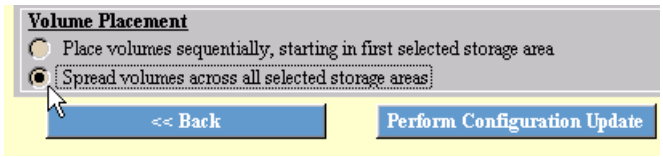

Figure 64. Volume placement

After all volumes to be defined are specified, the window looks like Figure 65. We can now perform the update by pressing Perform Configuration Update.

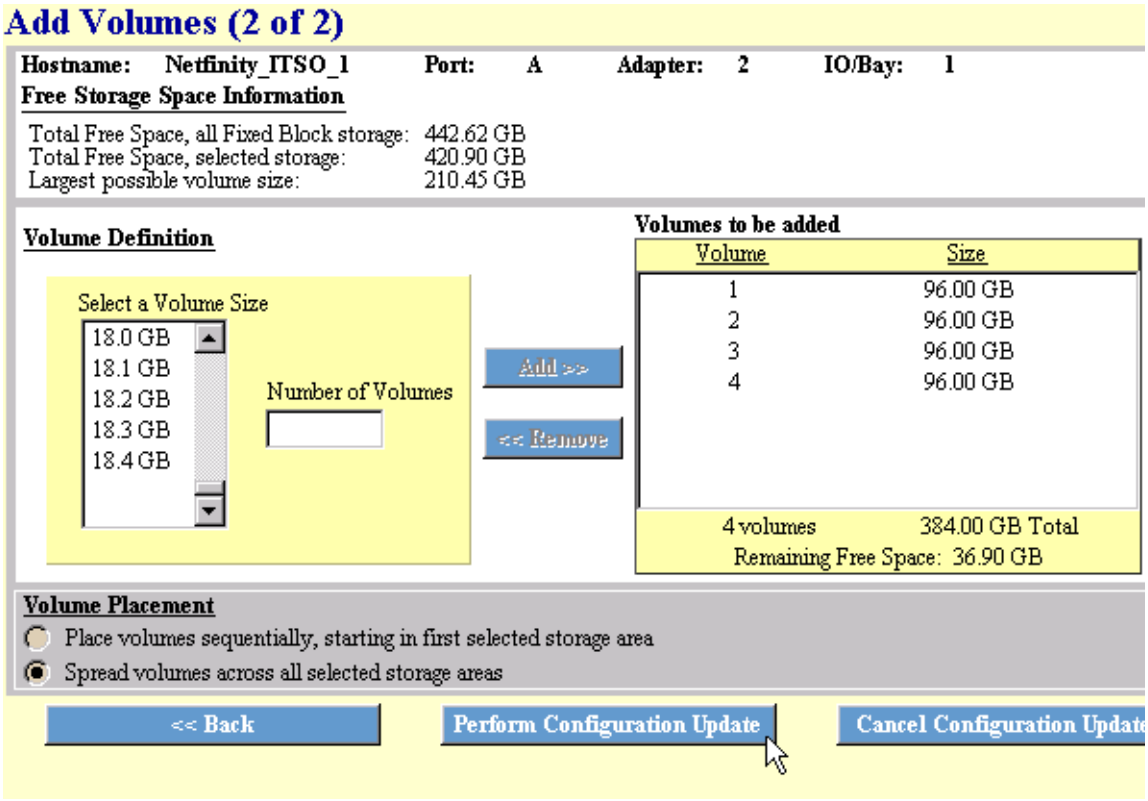

Figure 65. Performing the volume definitions

Again, we get a window warning us about a time intensive action and we are asked if we want to continue with the update. This is shown in Figure 66.

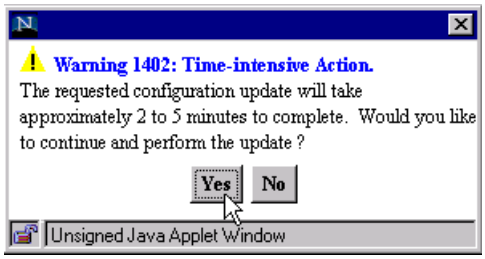

Figure 66. Warning window

As with every run of a configuration update, there is a progress window, as shown in Figure 67.

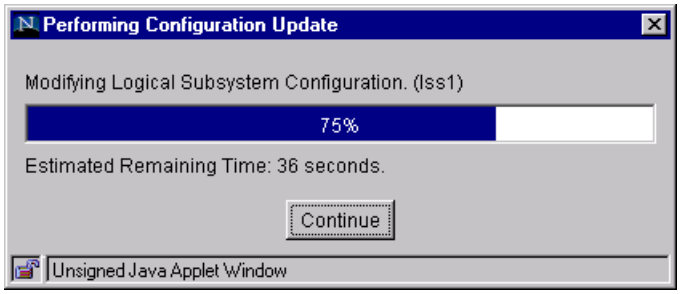

Figure 67. Progress window for volume definition

After the update has finished successfully, we are presented with the message shown in Figure 68.

Figure 68. Successful volume update

Pressing OK leads us back to the Adding Volumes (1 of 2) panel, from where we can start the process again and define other volumes.

To see what we have done, we press on the Storage Allocation button, shown in Figure 69. This is the same Storage Allocation button that we used in Figure 38 on page 76.

Figure 69. Storage Allocation button

This takes us back to the Storage Allocation -- Graphical View panel. There, we select the host Fibre Channel port that we configured, as shown in Figure 70.

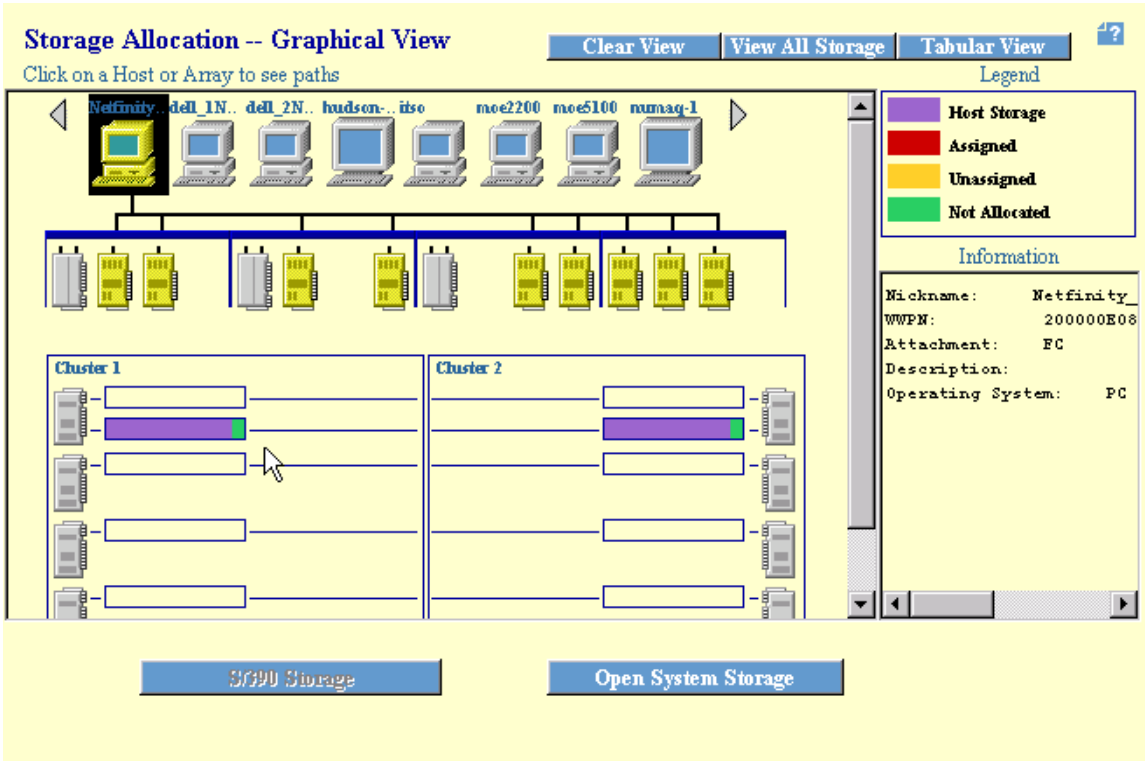

Figure 70. Storage allocation with host FC port selected

What this shows is the volumes defined for our host Fibre Channel port. Also, we see details of the selected host Fibre Channel port to the right of the window in the Information field.

Clicking on an ESS Fibre Channel port will draw lines to the volumes configured for the selected host Fibre Channel port, and the Information field now displays information of the configuration of the ESS Fibre Channel port. This is shown in Figure 71. This information is useful to see which ESS Fibre Channel port is configured for the appropriate topology. We show how to configure this in 4.1.7, “Configuring an ESS Fibre Channel port” on page 96.

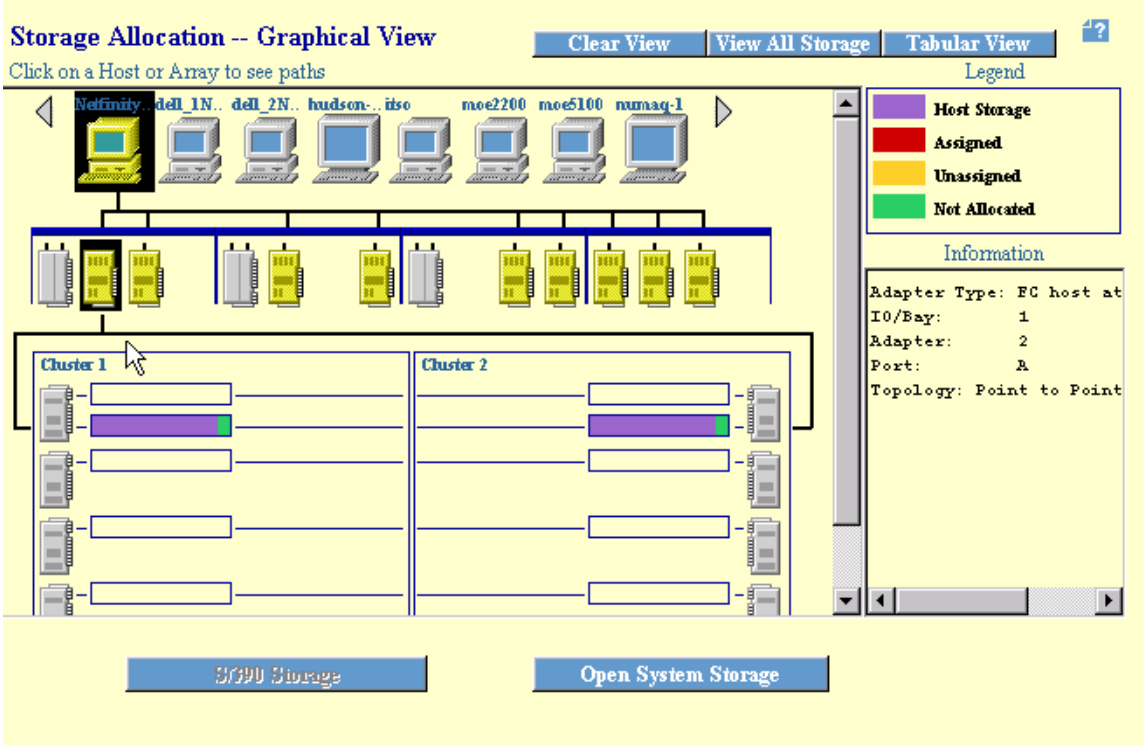

Figure 71. Storage allocation with host FC port and ESS FC port selected

To see the configured host Fibre Channel ports and the assigned volumes in the tabular view, we click on the Open System Storage button.

This takes us to the Open System Storage panel, as shown in Figure 72.

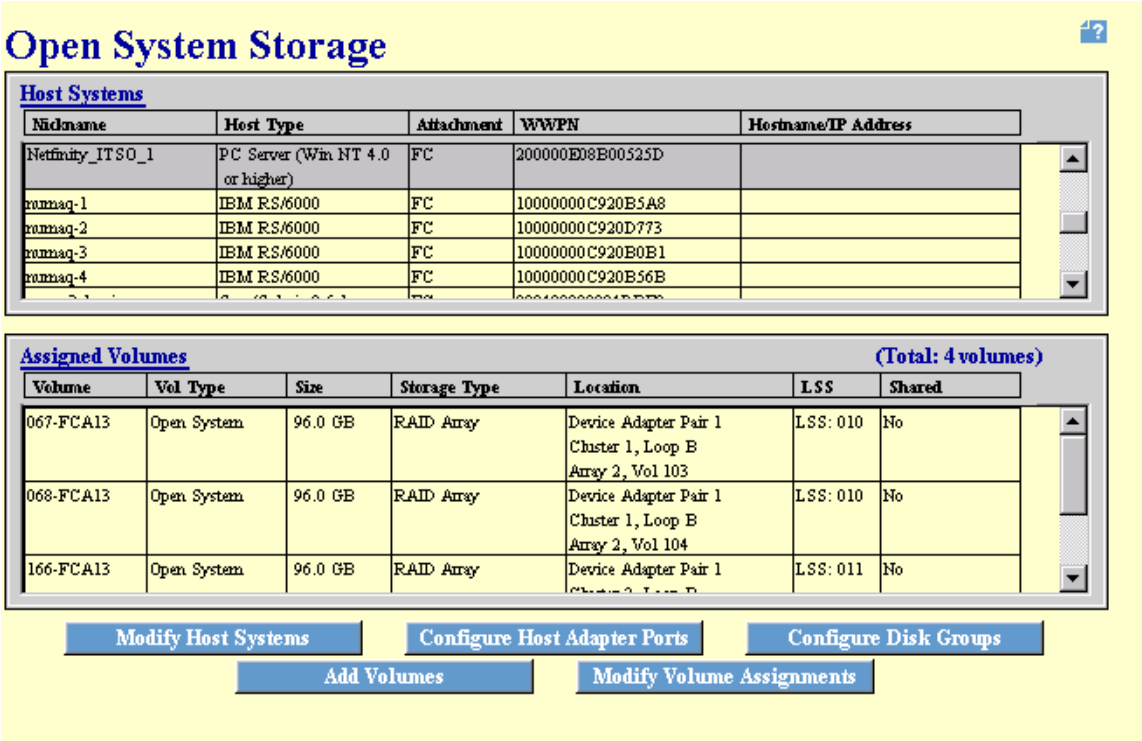

Figure 72. Host system port with associated volumes

Now when we select the host Fibre Channel port that we configured in the Host Systems field, the associated volumes appear in the Assigned Volume field with all the relevant details.

As a reminder of what we have done so far, we have:

- Configured a Fibre Channel port of our host system
- Configured disk groups
- Assigned volumes to the Fibre Channel host port (using the disk groups that we configured

Now, we will configure an ESS Fibre Channel port.

4.1.7 Configuring an ESS Fibre Channel port

Pressing the Configure Host Adapter Ports button takes us to the Configure Host Adapter Ports panel, as shown in Figure 73.

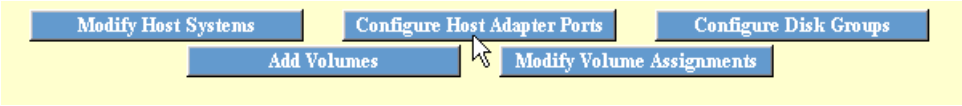

Figure 73. Entry point for configuring ESS FC ports

Again, the naming convention in the ESS specialist can be a little misleading. The ESS has built in Fibre Channel adapters, which, when used in a host system, would be host adapters. An ESS, in its simplest form, is a Fibre Channel storage device, not a host. We will refer to these as an ESS Fibre Channel adapter to avoid any confusion.

The Configure Host Adapter Ports panel is shown in Figure 74.

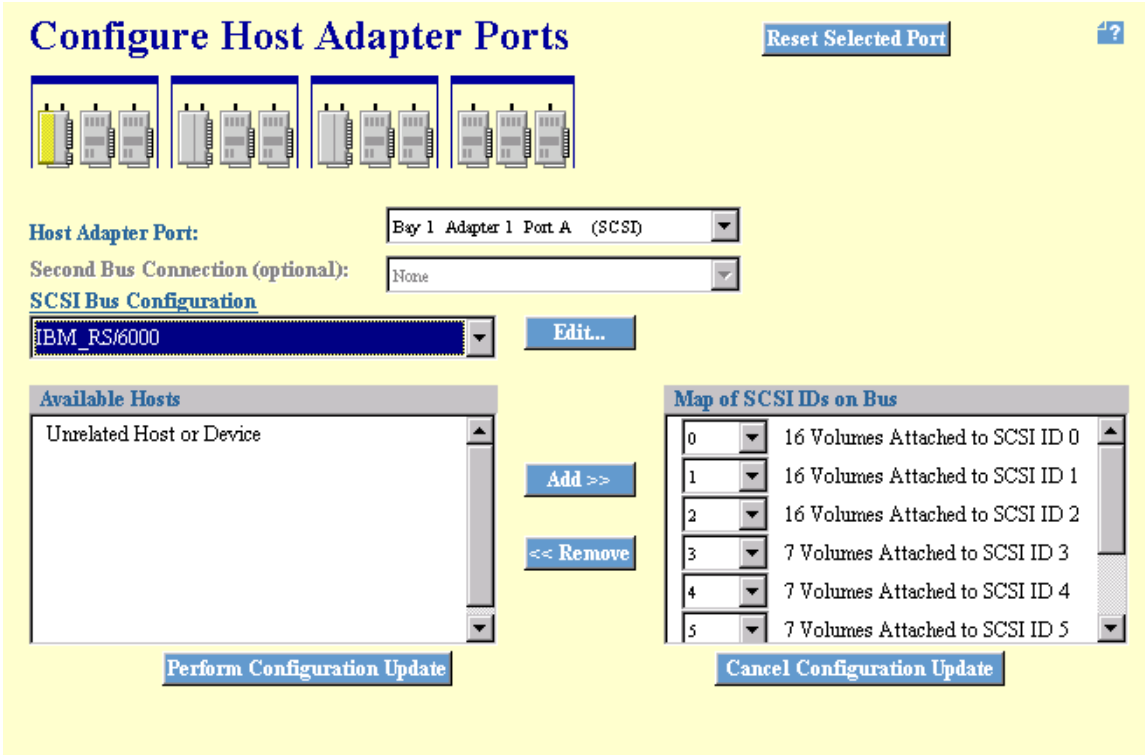

Figure 74. Configuring ESS interface ports - SCSI

4.1.7.1 Selecting the interface port

In this panel, we select the port of the ESS interface adapters to be configured depending on the type of the adapter. There are different configuration options for SCSI and Fibre Channel. For example, if we select a

SCSI port, we can configure the port for SCSI, as shown in Figure 74. However, we will focus on the FC part.

By selecting an FC port, we are presented with the Configure Host Adapter Ports panel, as shown in Figure 75.

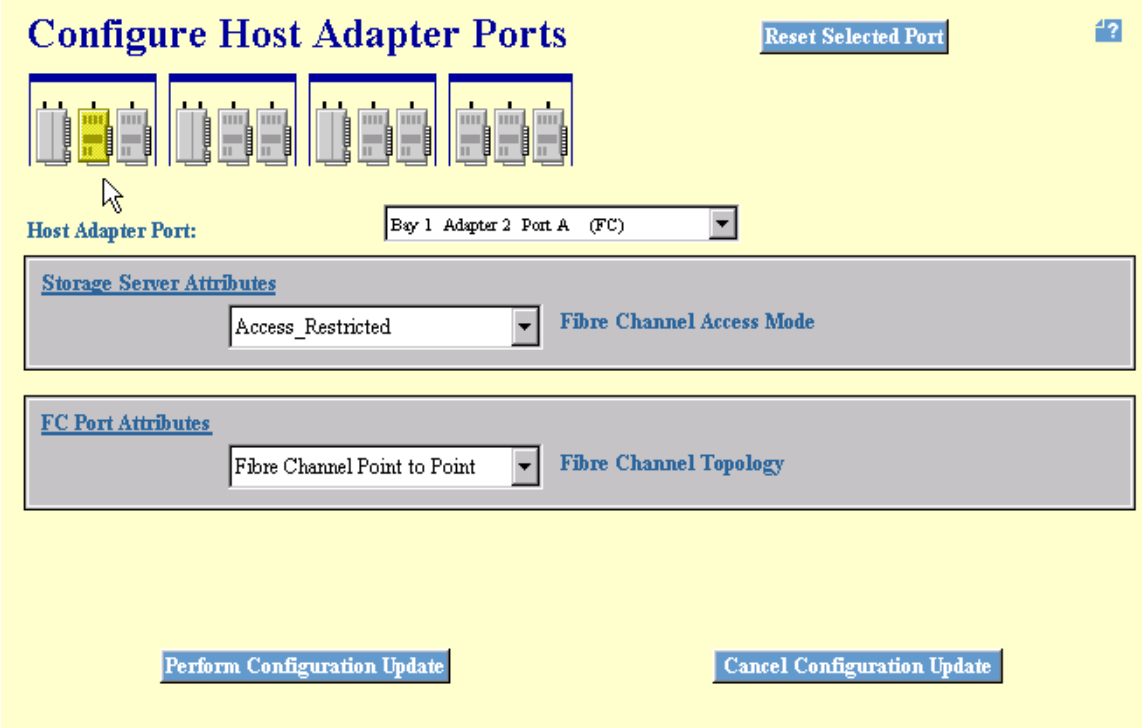

Figure 75. Configuring ESS interface ports - FC

The port we used is in bay 1, adapter 2, and is port A. It is called ‘A’ even it is the only port on the adapter.

4.1.7.2 Storage Server Attributes field

As shown in Figure 76, the Storage Server Attributes field — also known as the logical unit number (LUN) access mode — which specifies the channel access mode, can only be changed by an IBM service support representative (SSR).

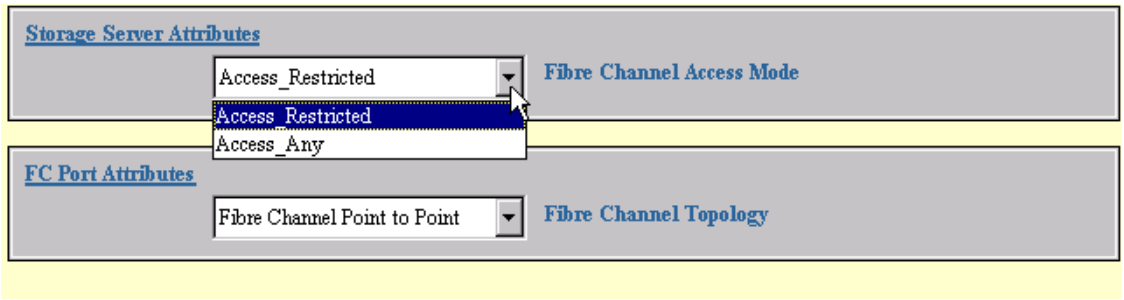

Figure 76. Storage Server Attributes field

There are two Fibre Channel Access Modes to which the ESS can be configured:

- Access_Any

  In Access_Any mode, any host system that is not defined in the ESS Specialist, can access all non-AS/400, open system logical volumes. In Access-Any mode, the ESS Specialist Web interface displays an anonymous pseudo-host, which is shown in Figure 77 and which is taken from another ESS, because the one we used here is configured for Access_Restricted. An anonymous host is not a real host system in disguise and connected to the storage server. It represents all Fibre Channel-attached host systems (if any) that are connected to the ESS and that are *not* configured on the ESS. The anonymous hosts do not have an access profile defined. You can access specific logical volumes that are defined in the ESS, by hosts that are not identified by the ESS.

- Access_Restricted

  In Access_Restricted mode, a connected host, which is not configured with the WWPN of its host FC adapter, cannot access any volumes. When this host gets configured by using the WWPN of its host FC adapter, it will only see the volumes for which it (the host FC adapter) is configured. Once a host is configured, there are no differences.

4.1.7.3 Access profiles

Whichever access mode is chosen, any Fibre Channel-attached host system that has an access profile can access only those volumes that are defined in the profile. Depending on the capability of the particular host system, an access profile can contain up to 256 or up to 4096 volumes.

The setup of an access profile is transparent to the user when you use the ESS Specialist Web user interface to configure the hosts and volumes in the ESS. Configuration actions that affect the access profile are:

- When you locate the WWPN for a Fibre Channel attached host system, you create the access profile for that host system. Initially the profile is empty. That is, it contains no volumes. The host cannot access any logical volumes that are already defined in the ESS.
- When you add new logical volumes to the ESS, the new volumes go to the host that you select.
- When you assign volumes to Fibre Channel attached hosts, you add to the access profile of the selected host system and to pre-existing volumes.
- When you remove a Fibre Channel attached host system from the ESS, you delete the host and its access profile.

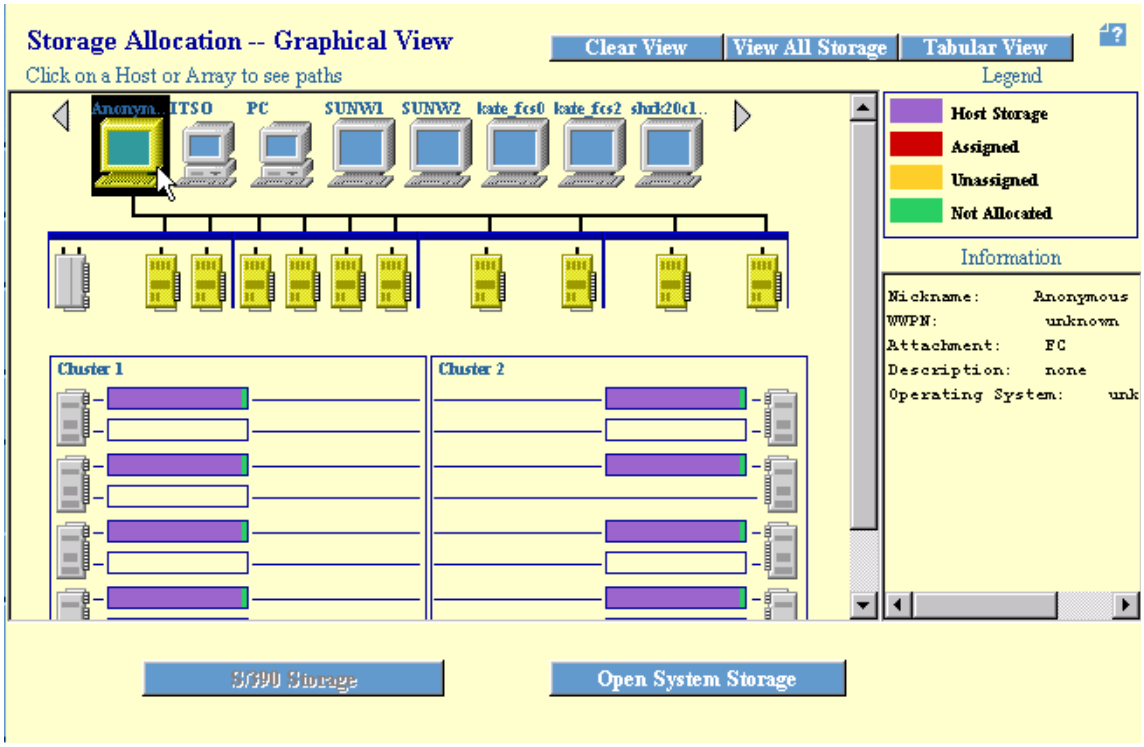

Figure 77. Anonymous host in access any mode

4.1.7.4 FC Port Attributes field

The second pull-down is for the Fibre Channel port attributes, and these are shown in Figure 78 and Figure 79. In Figure 78, the port selected is already configured for Fibre Channel Point to Point.

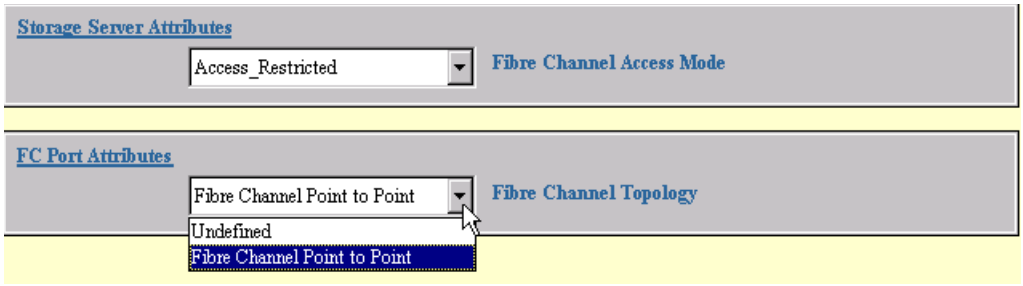

Figure 78. FC Port Attributes, configured for Point-to-Point

In Figure 79, the port is configured for Fibre Channel Arbitrated Loop.

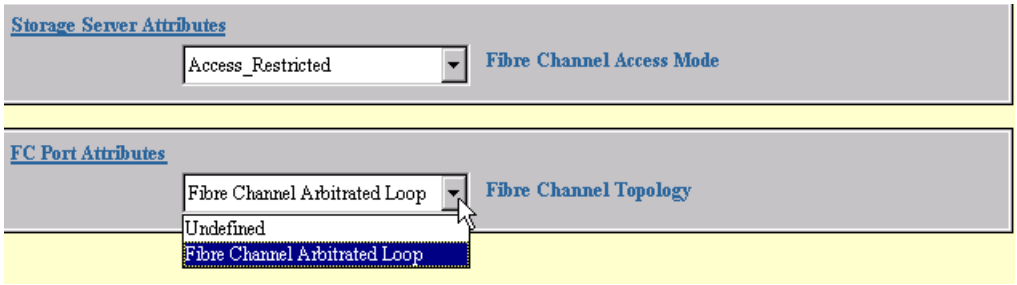

Figure 79. FC Port Attributes, configured for FC Arbitrated Loop

Obviously, it is here that we define the topology to which the port will be connected. The possible choices are Fibre Channel Point to Point and Fibre Channel Arbitrated Loop. We can choose between Fibre Channel Arbitrated Loop for FC-AL and Fibre Channel Point-to-Point for FC-PP *and* for FC-SW.

If a port is already configured for one topology and it has to be changed to another, the port must first be undefined. Then it is possible to choose the appropriate topology.

We want to undefine the Fibre Channel port in bay 4, adapter 3, and define this port to Fibre Channel Point to Point.

To undefine the port, we select it and highlight Undefined in the FC Port Attributes field. Once highlighted, we press the Perform Configuration Update button.

This will perform the update. As a result, the Fibre Channel port is returned to an undefined state, as shown in Figure 80.

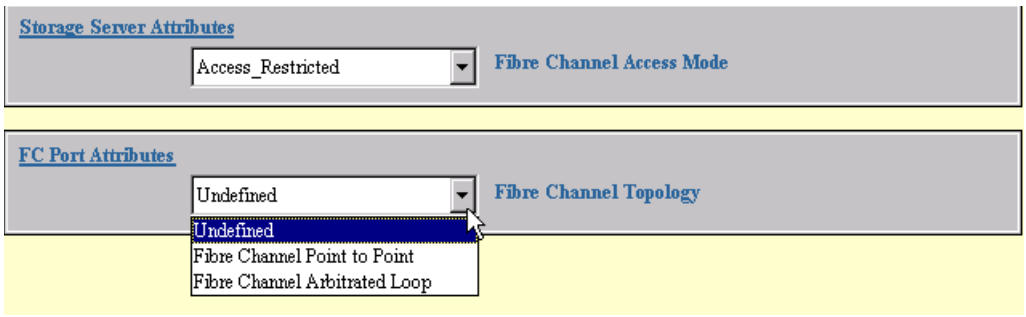

Figure 80. FC Port Attributes, undefined

Now we can choose between Fibre Channel Point to Point and Fibre Channel Arbitrated Loop. Again we need to press the Perform Configuration Update button

As we are in the process of performing the update, we get the progress window, shown in Figure 81.

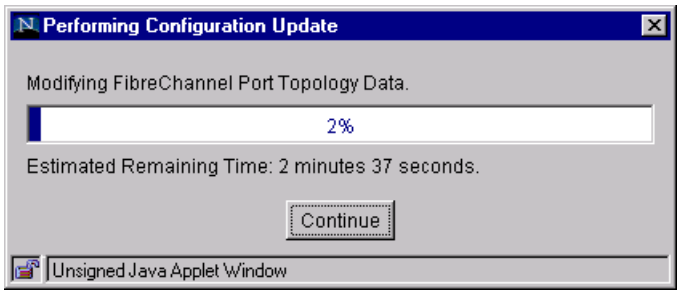

Figure 81. Progress window of changing the topology

Figure 82 shows a successful topology change of the port.

Figure 82. Successful topology change

Pressing the OK button takes us back to the Open System Storage panel, as shown in Figure 83.

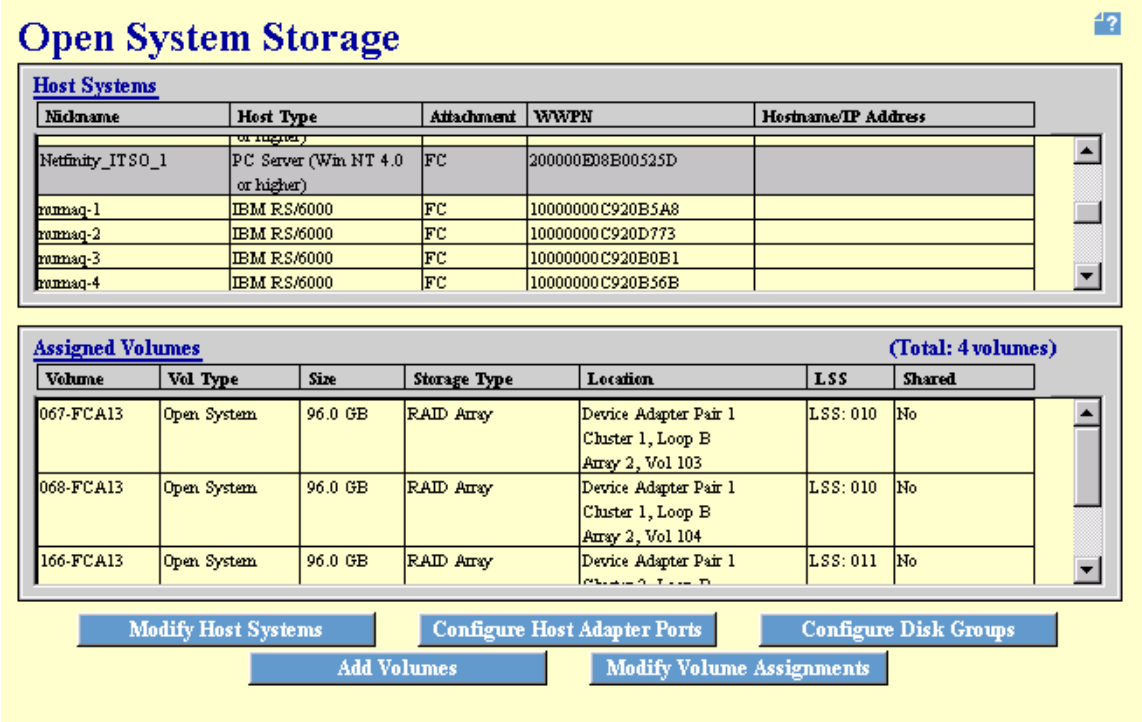

Figure 83. Open System Storage panel, no shared volumes

From here, we can continue to modify volume assignments. What we want to do is to share the volumes we defined with another host Fibre Channel port.

In the Shared column of the Assigned Volumes table in Figure 83, we can see that the volumes we created are not shared at this moment in time.

To do this, we first click on the Storage Allocation button to get a graphical representation of our ESS Fibre Channel port configuration. This is shown in Figure 84.

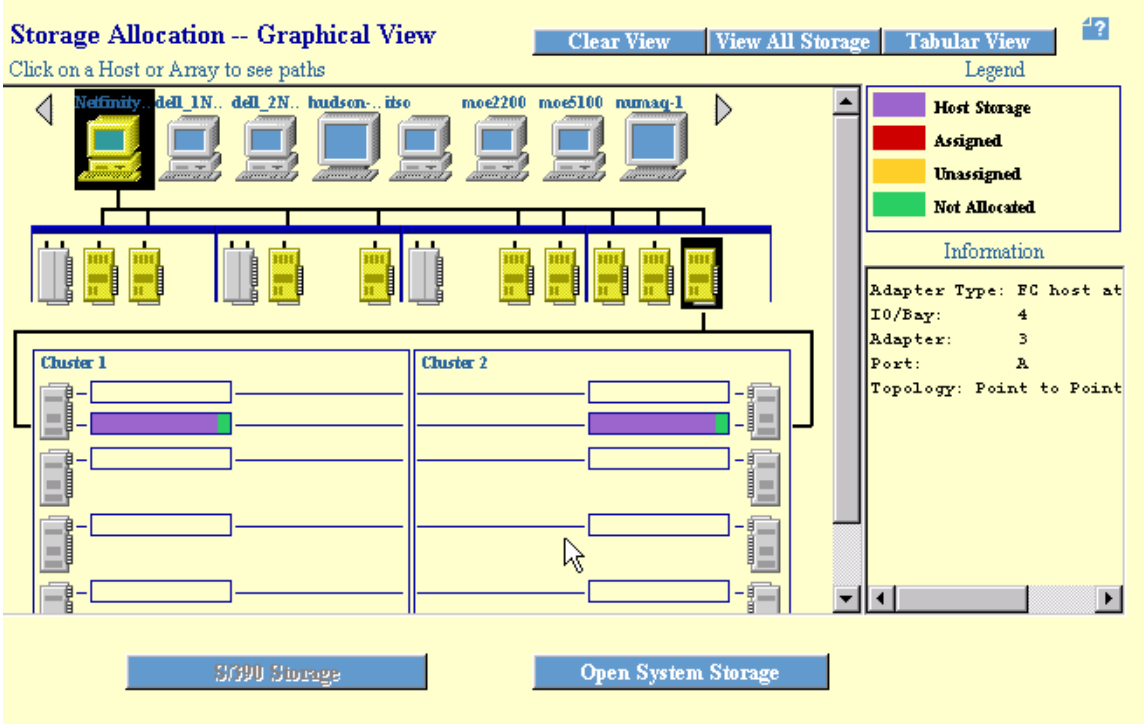

Figure 84. Selected host FC adapter and ESS FC adapter

In the Information field, we can verify that the port we selected to plug into the switched fabric is configured correctly.

4.1.8 Modifying volume assignments and sharing volumes

At this time, we have configured all that is necessary to enable the ESS to present volumes to a Fibre Channel host.

For further changes of the assignment of volumes to ports, we use the Modify Volume Assignments panel. For example, we can share volumes between host Fibre Channel ports. These host Fibre Channel ports can be in the same host which uses, for example, the Subsystem Device Driver (SDD), to increase the bandwidth and availability, or in different hosts, for example, for clustering.

We clicked on the Modify Volume Assignments button, as shown in Figure 85.

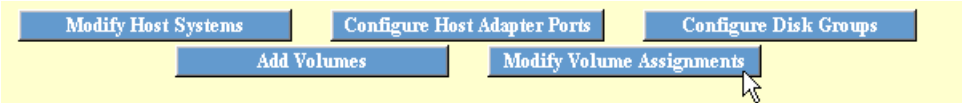

Figure 85. Entry point for modifying volume assignments

This takes us to the Modify Volume Assignments panel, as shown in Figure 86.

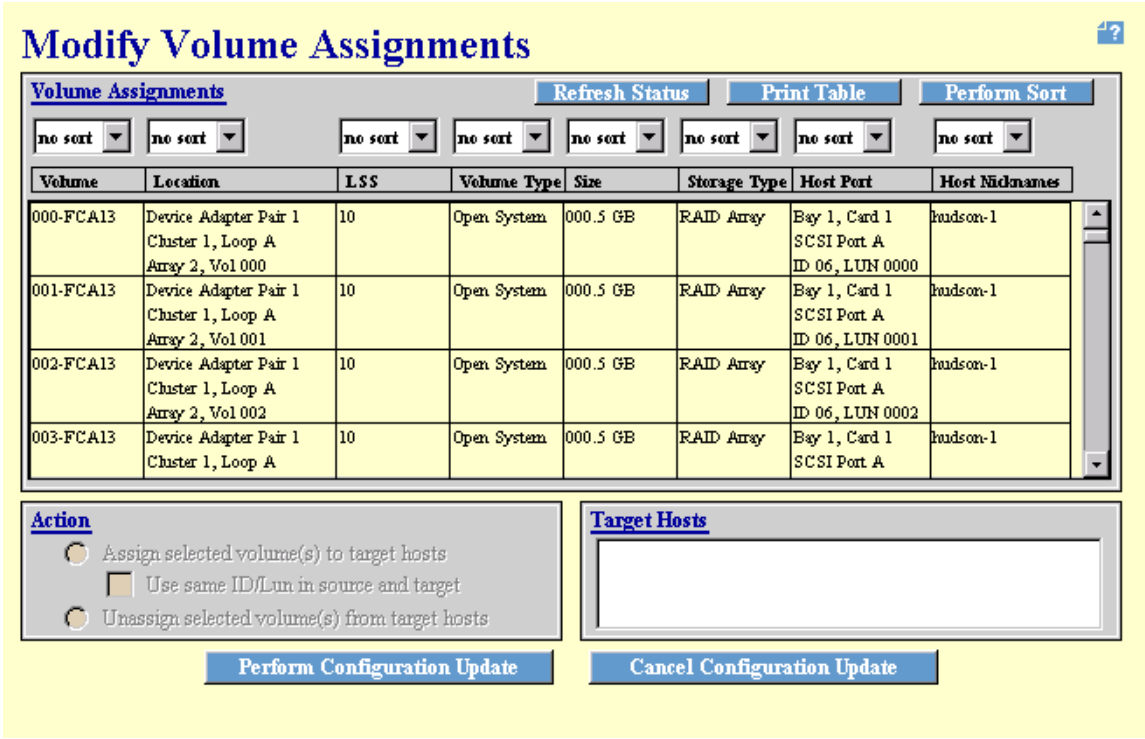

Figure 86. Modify Volume Assignments panel

The Volume Assignments table provides a comprehensive overview of all configured volumes with any associated information. We can also sort the table, specifying the column which we want to use as sort criteria.

There is an Action field, which is grayed out when no volume is selected, and a Target Hosts field. Now we show their interaction.

Figure 87 shows volumes selected and ready for modification.

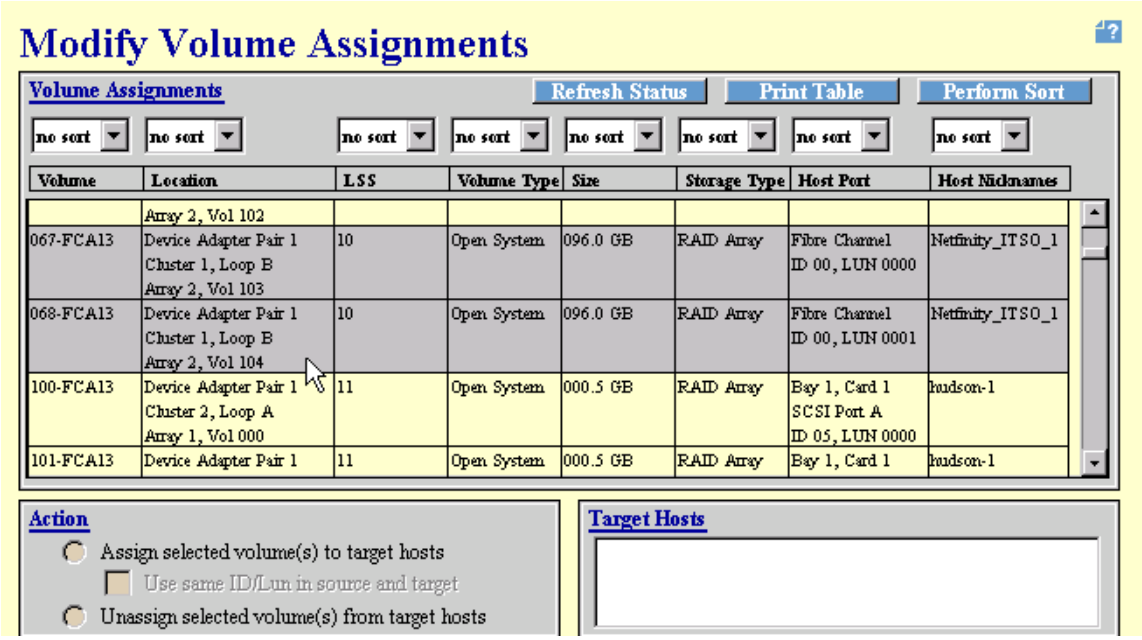

Figure 87. Selected volumes

The Action box is now accessible, as shown in Figure 88.

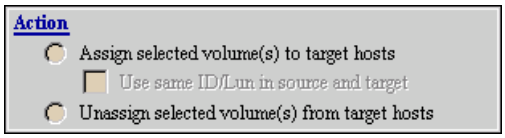

Figure 88. Accessible Action box

We Assign selected volume(s) to target hosts by selecting the radio button shown in Figure 89. From here, we can also Unassign selected volume(s) from target hosts.

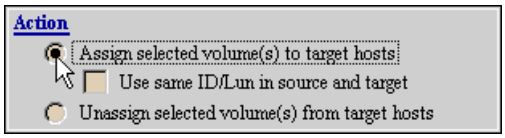

Figure 89. Checkbox to assign volumes

Selecting this allows us to choose the Target Hosts Fibre Channel port from those displayed in the Target Hosts field, as shown in Figure 90.

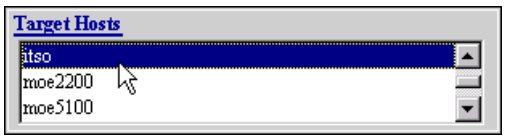

Figure 90. Field for target host Fibre Channel ports

We choose the host Fibre Channel port named ‘itso’ by highlighting its name. To apply the changes, we press the Perform Configuration Update button:

The progress bar shows the configuration update is taking place as shown in Figure 91.

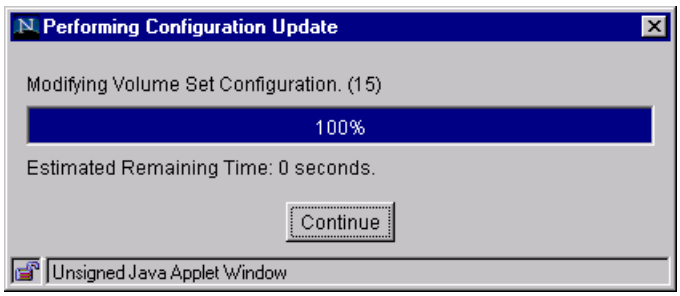

Figure 91. Applying volume assignment changes

If successful, we get the message shown in Figure 92.

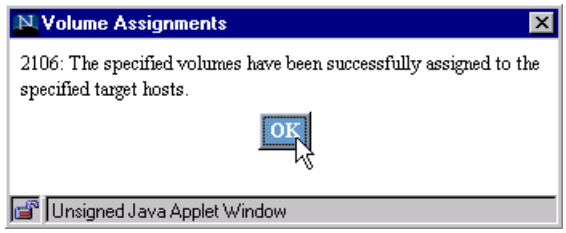

Figure 92. Volume Assignments successfully changed

So, if we press OK, we do not change back to the Open System Storage panel. Instead, we see the volume assigned to the host Fibre Channel port ‘itso’, as shown in Figure 93.

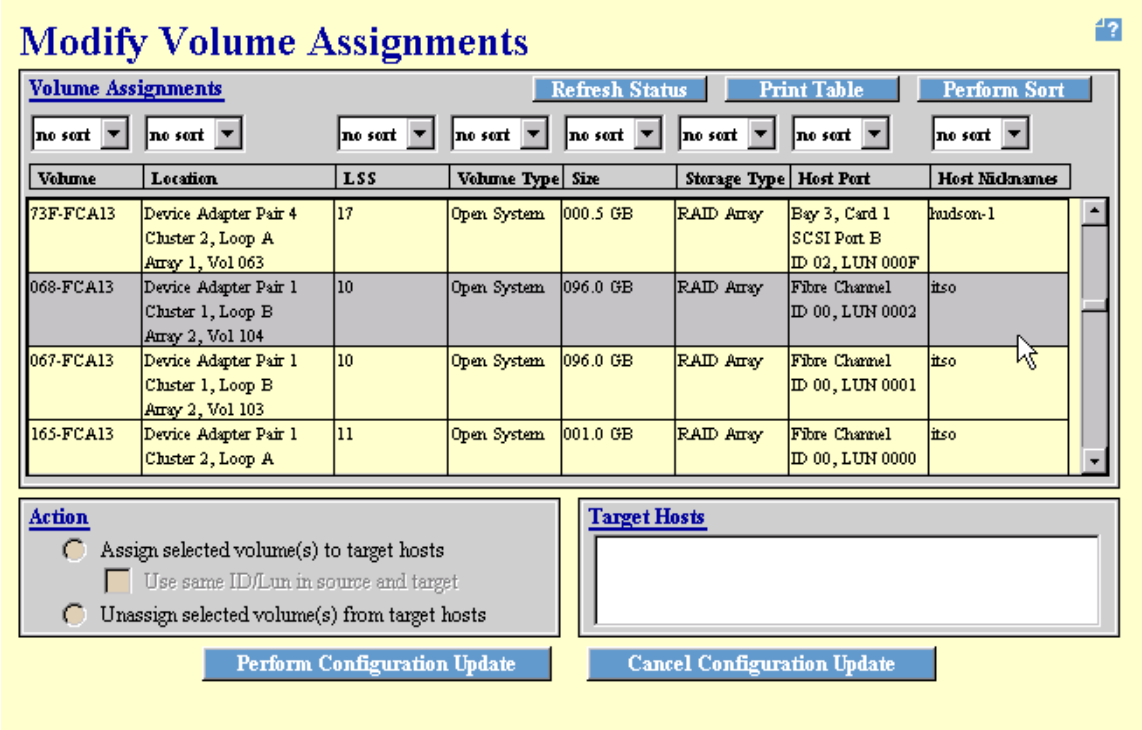

Figure 93. To 'itso' assigned volume

The same volume is also assigned to the host Fibre Channel port of our other Fibre Channel host Netfinity_ITSO_1, as shown in Figure 94.

068-FCA13 Formatting (84%)	Device Adapter Pair 1 Cluster 1, Loop B Array 2, Vol 104	10	Open System	096.0 GB	RAID Array	Fibre Channel ID 00, LUN 0001	Netfinity_ITSO_1
068-FCA13 Formatting (84%)	Device Adapter Pair 1 Cluster 1, Loop B Array 2, Vol 104	10	Open System	096.0 GB	RAID Array	Fibre Channel ID 00, LUN 0002	itso
001-FCA13	Device Adapter Pair 1 Cluster 1, Loop A Array 2, Vol 001	10	Open System	000.5 GB	RAID Array	Bay 1, Card 1 SCSI Port A ID 06, LUN 0001	hudson-1
038-FCA13	Device Adapter Pair 1 Cluster 1, Loop A	10	Open System	000.5 GB	RAID Array	Fibre Channel ID 00, LUN 0052	moe5100

Figure 94. To 'Netfinity_ITSO_1' assigned volume

However, if we use the back button of our browser, we return to the Open System Storage panel, as shown in Figure 95.

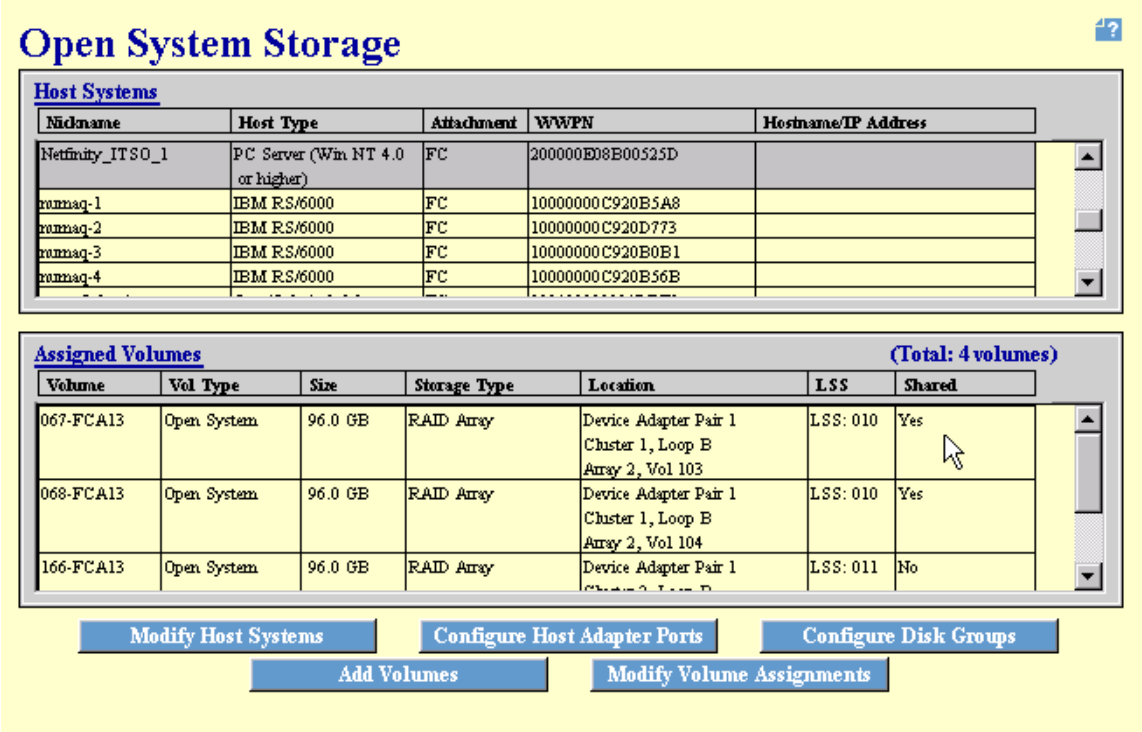

Figure 95. Open System Storage panel, with shared volumes

If we look at the Shared field, we can see that this volume is now shared.

4.2 Related information

For more information on the ESS, see:

IBM Enterprise Storage Server Introduction and Planning Guide, 2105 Models E10, E20, F10 and F20, GC26-7294

- Introduces the product and lists the features you can order. It also provides guidelines on planning for the installation and configuration of the ESS.

IBM Enterprise Storage Server User's Guide, 2105 Models E10, E20, F10 and F20, SC26-7295

- Provides instructions for setting up and operating the ESS.

IBM Enterprise Storage Server Host Systems Attachment Guide, 2105 Models E10, E20, F10 and F20, SC26-7296

- Provides guidelines for attaching the ESS to your host system.

IBM Enterprise Storage Server SCSI Command Reference, 2105 Models E10, E20, F10 and F20, SC26-7297

- Describes the functions of the ESS. It provides reference information for UNIX and AS/400 hosts, such as channel commands, sense bytes, and error recovery procedures.

IBM Enterprise Storage Server System/390 Command Reference, 2105 Models E10, E20, F10 and F20, SC26-7298

- Describes the functions of the ESS and provides reference information for System/390 hosts, such as channel commands, sense bytes, and error recovery procedures.

ESS Web Interface User's Guide for ESS Specialist and ESS Copy Services, SC26-7346

- Provides instructions for using the IBM StorWatch ™ Enterprise Storage Server Web interface, ESS Specialist.

IBM Storage Solutions Safety Notices, GC26-7229

- Provides translations of the danger notices and caution notices that IBM uses in ESS publications.

IBM Enterprise Storage Server Configuration Planner, SC26-7353

- Provides work sheets for planning the logical configuration of the ESS. This book is not available in hard copy. This book is only available on the following Web site:

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

IBM Enterprise Storage Server Quick Configuration Guide, SC26-7354

- Provides flow charts for using the StorWatch Enterprise Storage Server Specialist. The flow charts provide a high-level view of the tasks the IBM service support representative performs during initial logical configuration. You can also use the flow charts for tasks that you might perform during modification of the logical configuration.The hard copy of this booklet is a 9-inch by 4-inch fanfold.

Chapter 5. Implementing an IBM managed and unmanaged hub

In this chapter we show how an unmanaged hub and a managed hub can be implemented. For the managed hub we also show how to use a feature called QuickLoop and QuickLoop zoning.

The products we describe are the:

- IBM Fibre Channel Storage Hub, 2103-H07
- IBM SAN Fibre Channel Managed Hub, 35341RU

These are seen as entry level components of a SAN.

5.1 ITSO environment

This setup, as shown in Figure 96, is mainly a low cost, entry level solution for connecting one or more storage devices, to one or more servers. It is not a very scalable solution, and should not be chosen if many devices are to be connected later. A hub can also be used to connect to a remote location to extend the distance.

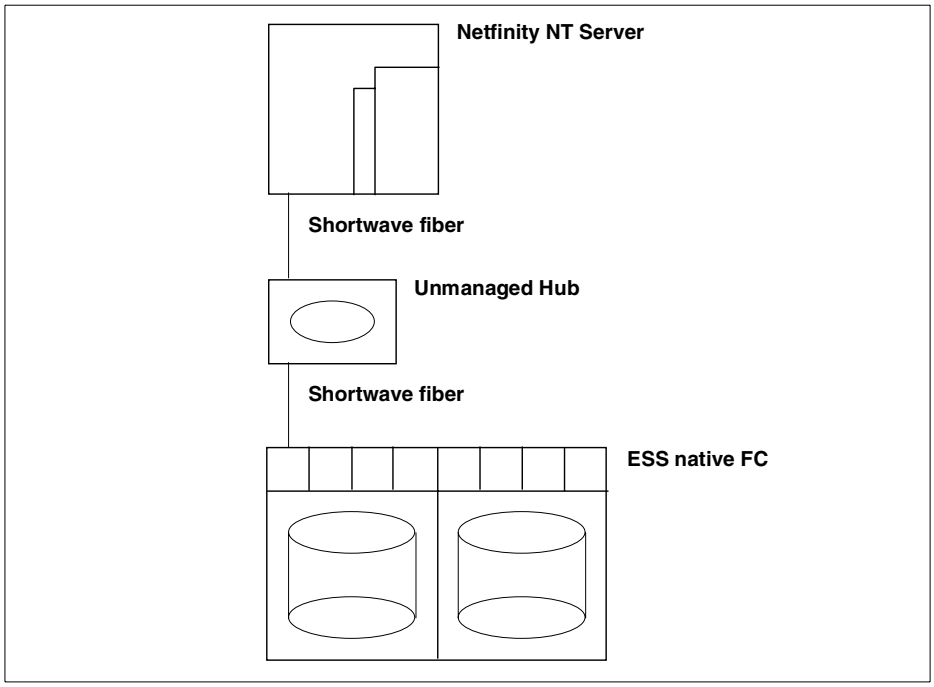

Figure 96. Simple loop setup

5.1.1 Hardware overview

Here is the hardware setup:

- 1 IBM Fibre Channel Storage Hub, 2103-H07
  - 7 ports
  - 2 shortwave gigabit interface converters (GBICs), 2010
  - 2 shortwave cables, 9701 5m, 9702 25m
- 1 Netfinity NT Server 5500
  - 1 Fibre Channel (FC) HBA QLogic 2100, FC_AL
- 1 IBM ESS 2105-F20 with Feature Code 3022 (native FC adapter)
  - 9 FC adapters, 9 ports
  - 3 SCSI adapters, 6 ports

5.1.2 Software overview

Here is the software setup:

- Netfinity: Microsoft Windows NT 4.0 Service Pack 5
- QLogic 2100
- ESS: Microcode (1.1)

5.1.3 Configuring the ESS

To configure the ESS, refer to Chapter 4, “Configuring the ESS with native Fibre Channel” on page 63.

5.1.4 Configuring the host bus adapters

To configure the QLogic 2100, refer to 3.1.1, “Installing the QLogic 2100F adapter card” on page 53.

5.2 Configuring the unmanaged hub

Our first SAN scenario consists of an IBM Enterprise Storage Server 2105-F20 (IBM ESS F20) with native Fibre Channel attachment, linked to an IBM Netfinity 5500 server through an IBM 2103-H07 Fibre Channel Storage Hub. Our planned configuration is shown in Figure 97.

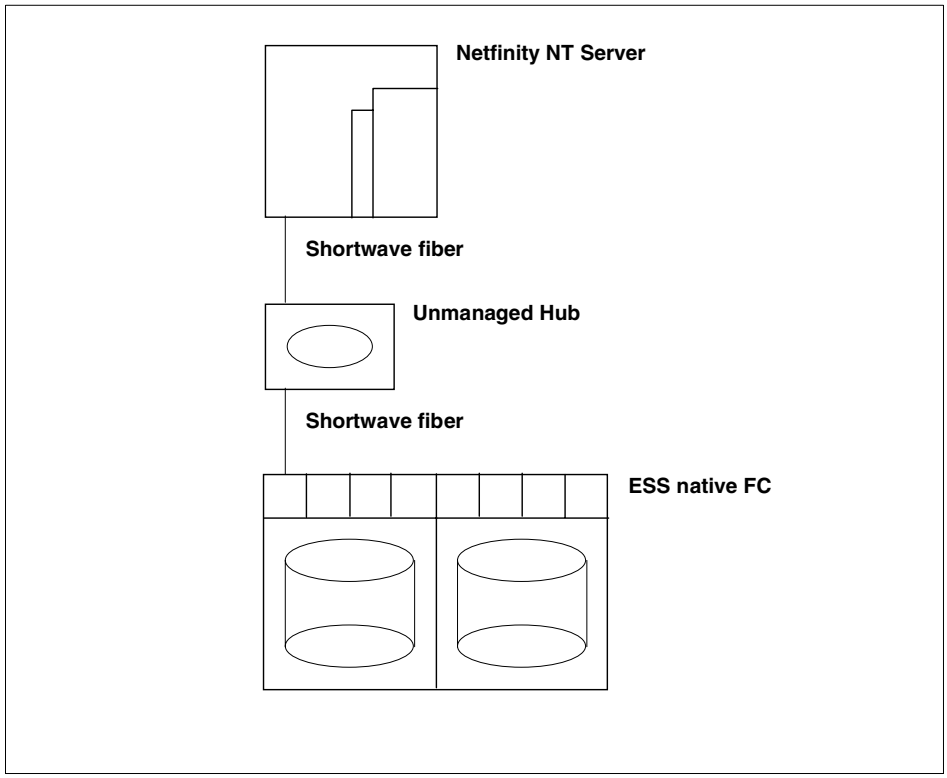

Figure 97. Simple SAN with hub

5.2.1 IBM Fibre Channel Storage Hub

Fibre Channel Storage Hubs are seen as entry level components in SAN fabric installations.

The Fibre Channel Storage Hub is designed to provide a centralized point of connectivity, to provide loop fault tolerance, and to simplify configuration management.

Fibre Channel products that are commonly interconnected to the Fibre Channel Hub are Fibre Channel host bus adapters, FC-AL storage devices, and FC-AL storage arrays.

IBM today offers two kind of hub solutions:

- The unmanaged IBM Fibre Channel Storage Hub 2103-H07
- The managed IBM Fibre Channel Storage Hub 3534-1RU

5.2.2 IBM Fibre Channel Storage Hub, 2103-H07

The IBM Fibre Channel storage Hub 2103-H07 used in our first scenario, is a 7-port central interconnection for Fibre Channel Arbitrated Loops that follow the ANSI FC-AL standard. In Figure 98 we show the hub.

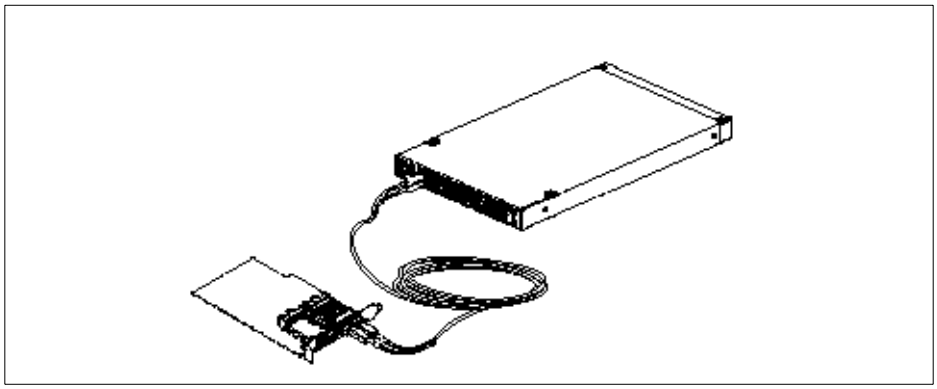

Figure 98. Fibre Channel Hub and Interface Connector

5.2.3 Hub configuration

The IBM Fibre Channel Storage Hub interconnects multiple servers and storage systems, over fiber-optic media, and transfers data at speeds up to 100 MB/s.

Each port requires a gigabit interface converter to connect it to each attached node. The Fibre Channel Storage Hub supports any combination of shortwave or longwave optical GBICs. We show a GBIC in Figure 99.

Figure 99. Gigabit Interface Converter

The GBICs are hot-pluggable into the IBM Fibre Channel Storage Hub, which means you can add host computers, servers, and storage modules 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 port, that port is automatically bypassed. The remaining hub ports continue to operate normally with no degradation of system performance. Conversely, if you plug a GBIC into the Fibre Channel Storage Hub, it will automatically be inserted and become a node on the loop, if valid Fibre Channel data is received from the device.

5.2.4 Hub installation

In Figure 100 we show the front panel of the IBM Fibre Channel Storage Hub, 2103-H07.

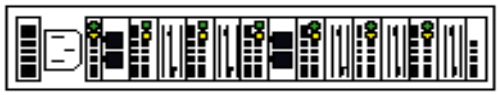

Figure 100. FC Storage Hub 2103-H07 front panel

The picture shows the front panel with the seven slots for GBICs. After GBIC installation they represent the FC Storage Hub ports 0 to 6.

5.2.4.1 GBIC insertion

The unmanaged hub is as close to plug and play as you are likely to encounter in a SAN installation. Of course, there is more to consider, but in our installation, we will show how to install GBICs for those that are new to the Fibre Channel world. To insert the GBIC follow these steps:

1. Remove the plastic cover (1) from the GBIC (2). This is shown in Figure 101.

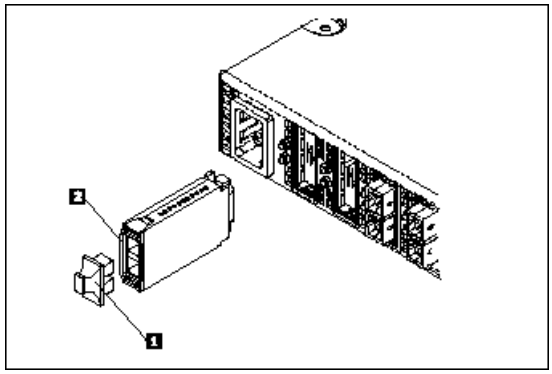

Figure 101. Insert GBIC

2. Slide the GBIC (2) into the port.
3. Connect the fiber-optic cable (3) into the installed GBIC.

   The GBIC housing has an integral guide key that is designed to prevent improper insertion, and this is shown in Figure 102.

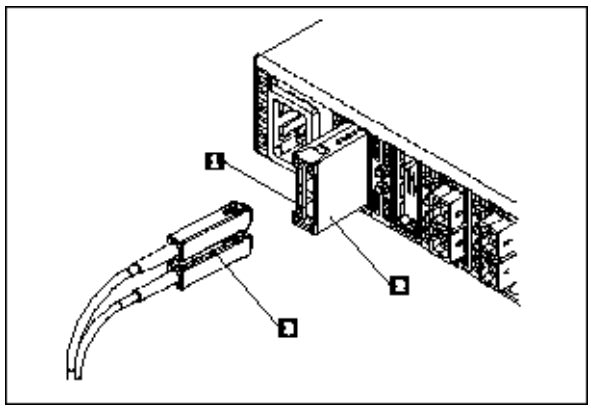

Figure 102. Insert cable into GBIC

4. Once you have installed the GBICs in your hub, attach the hub ports, using standard 50/125 or 62.5/125 FC cable, to your Fibre Channel host and storage device.

The FC-AL was now operational and we powered on the hub. The Netfinity server and the ESS were already connected and running.

By looking at the Device Active LEDs on the FC Storage Hub ports, as shown in Figure 103, and ensuring that they were green, we knew that all the connections were working.

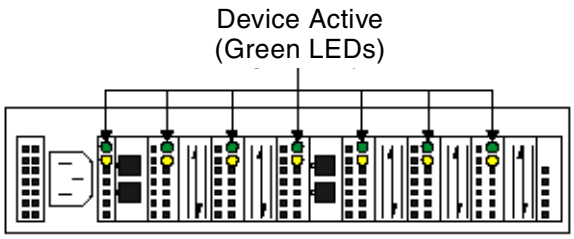

Figure 103. Device Activity LEDs in FC Storage Hub

Now we checked to see whether we can see the same ESS logical volumes as though we were directly attached to the ESS.

5.2.5 Checking disk access

We used the Windows NT Disk Administrator to view the available disks, or ESS LUNs.

In Figure 104, we show the Disk Administrator window from the Windows NT system.

The system already had 44 disks available to it, but the 1 GB volume we assigned to it from the ESS was not recognized at this stage.

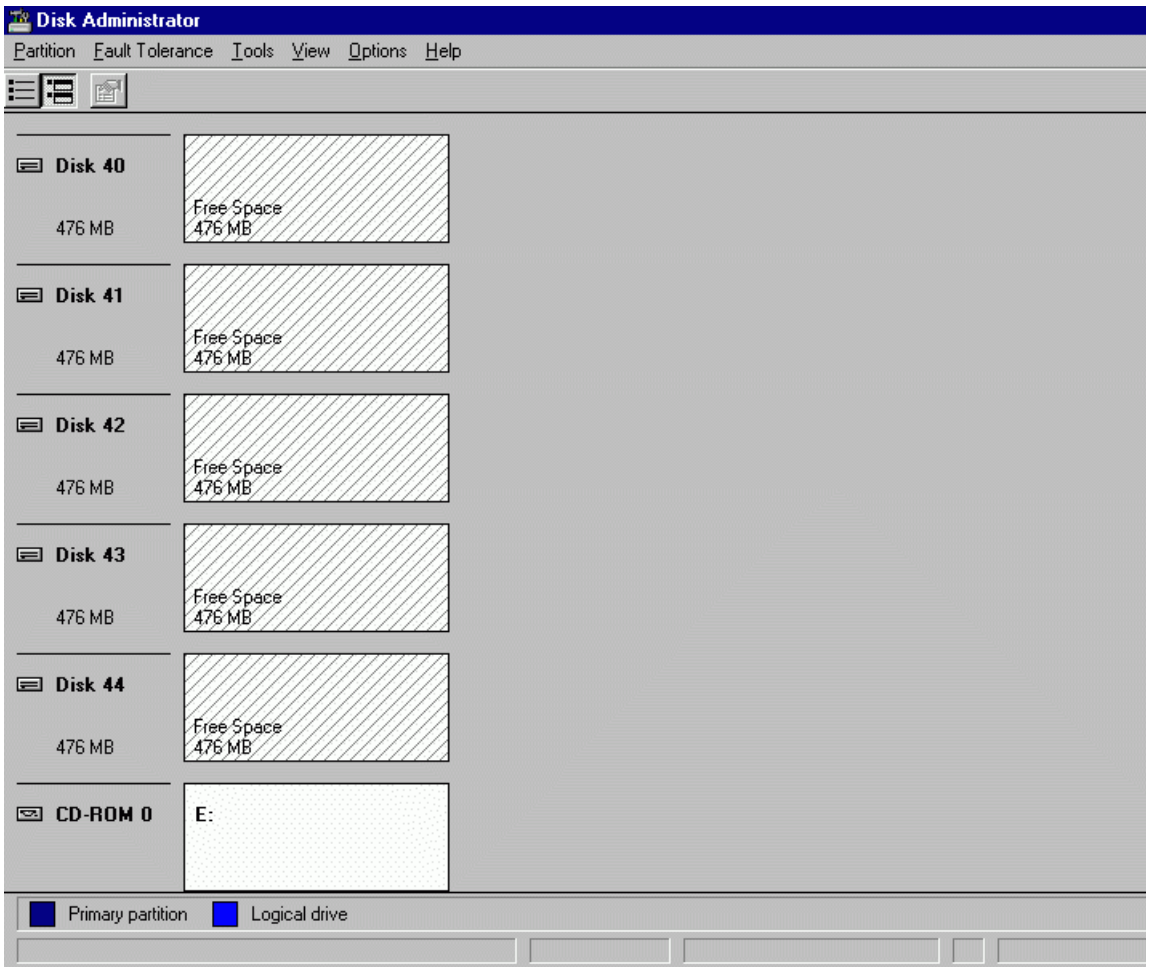

Figure 104. Disk Administrator view before reboot

We rebooted the Windows NT system to let the Windows operating system rescan for attached devices. Once this was done, an extra disk, Disk 45, now became available and ready for use.

In Figure 105, we show that Disk 45, with a 953 MB capacity, is now ready to be assigned and formatted for use as any regular disk.

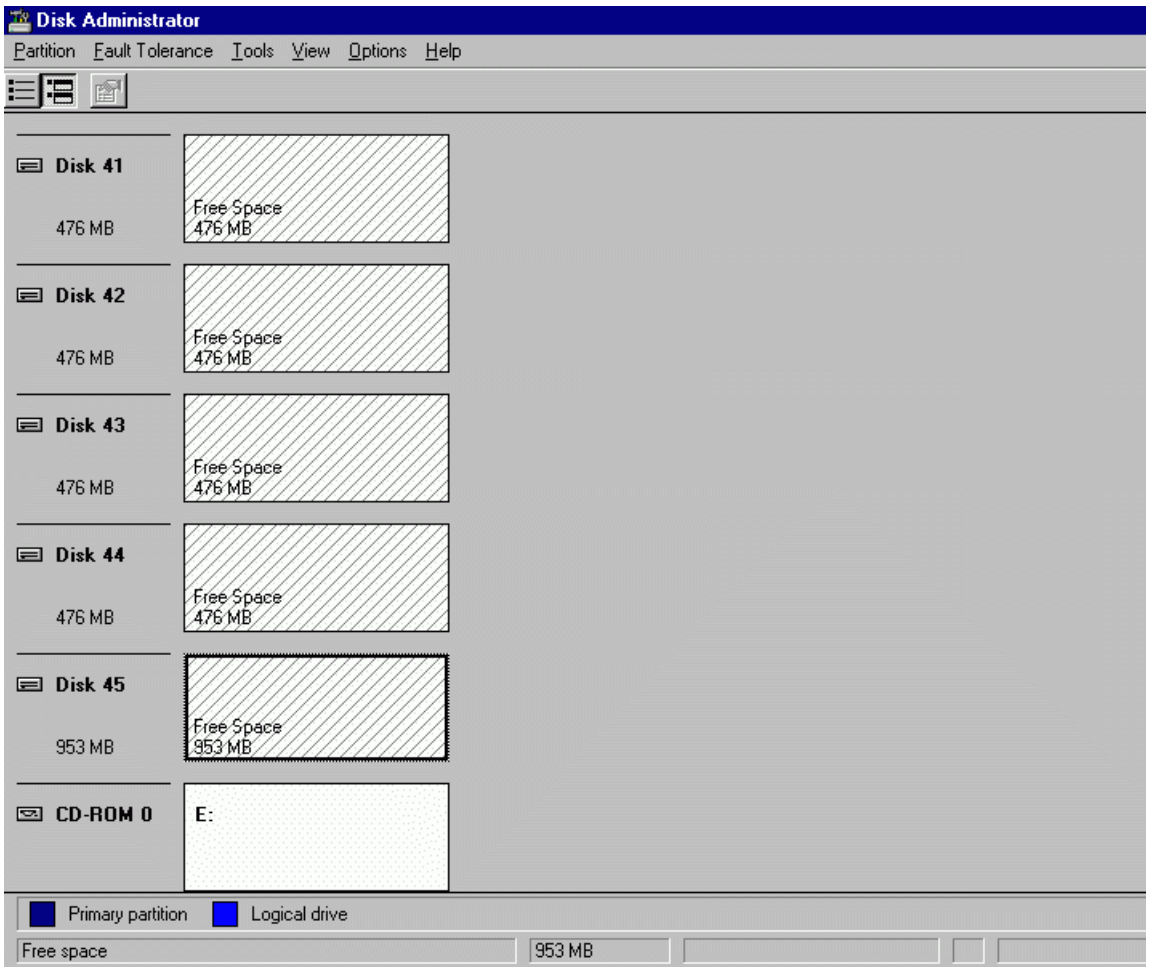

Figure 105. Disk Administrator view after reboot

A second Windows NT host was also attached to the hub, so that there were two analogous host operating systems in a single loop. We assigned the same volume in the ESS, that was assigned to the first host, to this second host, and then rebooted the second host system. It recognized the new volume and was able to access it.

To ensure data integrity, clustering, or shared access, software must be installed in both hosts.

Note

Two hosts, or initiators, attached to a hub in a single loop to an ESS is not supported by IBM.

Although the two host example is fully working, because of the loop initialization process, we do not recommend it. What we recommend is to have a second host with clustering software like the Microsoft Cluster Server, in a loop to take over operation if one host fails. However, this is still not a high availability storage configuration, because the loop itself provides no redundancy, and is therefore, potentially, a single point of failure. This problem could be solved by using two loops.

In terms of scalability of bandwidth, one FC-AL loop by itself is not scalable. All devices share the bandwidth of 100 MB/s, rather than that offered by the managed hub.

5.3 IBM SAN Fibre Channel Managed Hub

The IBM SAN Fibre Channel Managed Hub, 35341RU, is an 8-port fibre channel hub that consists of a system board with connectors for supporting up to eight ports. This includes seven fixed, short wavelength ports, one pluggable GBIC port, and an operating system for building and managing a switched loop architecture.

The hub is supported on IBM PC, Netfinity servers, and other Intel-based servers.

The latest support matrix, including adapters and operating system requirements, can be found at the following Web page:

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

In Figure 106, we show the faceplate of the IBM SAN Fibre Channel Managed Hub, 35341RU. The ports are numbered sequentially starting with zero for the left most port.

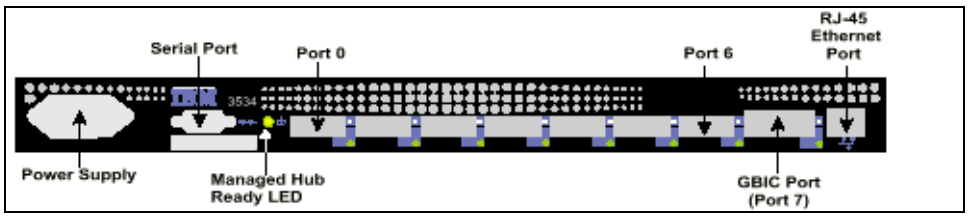

Figure 106. Managed hub

The system board is enclosed in an air-cooled chassis, which may be either mounted in a standard rack or used as a stand-alone unit.

The chassis includes a power supply, an RJ-45 Ethernet connection for 3534 Managed Hub set up and management, and a serial port. If the default address is not known, the serial port is used for recovering the factory settings and initial configuration of the IP address.

The hub can accommodate one GBIC module and can be connected to one other managed hub, to expand the loop capabilities to 14 ports. It can also be connected with a single port into a SAN fabric as a loop extension.

The hub may be managed using the serial port or the 10/100BaseT Ethernet port. Management interfaces include Telnet or Web-based management using the IBM StorWatch SAN Fibre Channel Managed Hub Specialist. This is similar to the StorWatch Fibre Channel Switch Specialist.

All network cable connection is through 50/125 or 62.5/125 Short Wave Length (SWL) cable or through 9/125 Long Wave Length (LWL) cable.

Fibre Channel Storage Hubs are designed to provide a centralized point of connectivity, to provide loop fault tolerance, and to simplify configuration management. Specifically designed for entry-level workgroup FC-AL applications, the hubs provide considerable flexibility in configuring loops and segmenting them for performance or high-profile availability applications.

5.3.1 The ITSO environment

From our ITSO environment, we show how to implement the IBM SAN Fibre Channel Managed Hub, 35341RU.

5.3.2 The hardware involved

We used this hardware in our configuration:

- 1 IBM PC with 1 FC HBA from QLogic

- 1 Fibre Channel RAID Storage server IBM 2102-F10 with native FC adapter
- 1 IBM 3534 managed hub
  - 7 fixed optic ports
  - 0 hot pluggable GBIC port (empty)

5.3.3 The software involved

We used this software in our configuration:

- Microsoft Windows NT 4.0 Service pack 5
- IBM 3534 microcode level 2.1.3

5.4 Installing the IBM SAN Fibre Channel Managed Hub, 35341RU

We recommend that you use a pre-installation checklist.

An example of a pre-installation checklist is detailed in the *IBM SAN Fibre Channel Managed Hub 3534 Service Guide,* SY27-7616 and the *IBM SAN Fibre Channel Managed Hub 3534 User's Guide,* GC26-7391.

The checklist ensures a successful installation and includes checks on the host operating system, host bus adapter, storage devices, cables and network parameters, the most important of which is the hub IP address.

5.4.1 Setting the IP address

The 3534 Managed Hub is shipped from the factory with a default IP address (10.77.77.77). This IP address is printed on the label on the top front edge of the 3534 Managed Hub. This address is for the external Ethernet connection.

If you can, use this default address to attach to your local area network to establish a network connection to the hub. In Figure 107 we show the Ethernet and serial port locations.

You can change this IP address later using a Telnet command, or by using the StorWatch Specialist, issued from any server having access to the same LAN. This is the easiest way to set the IP address.

Your system or network administrator will advise if the default address can be used.

If using the default IP address is not possible, the IP address will have to be set using the Ethernet port or the serial port. Set the IP address using the

information provided by the system administrator and record this on the pre-installation checklist.

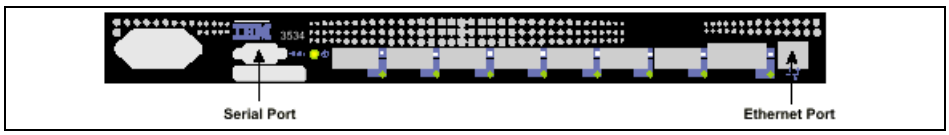

Figure 107. IBM 3534 management ports

For a comprehensive description of how to set the IP address, using either the serial, or the Ethernet port, refer to *IBM SAN Fibre Channel Managed Hub 3534 Service Guide,* SY27-7616.

The hub we used in our test scenario had been installed before, and we had to change the IP address.

In Figure 108, we show the front panel of the hub. Port 5 is operating, no GBIC is installed, and the left fan has stopped working. This is indicated by a red color and no movement of the fan.

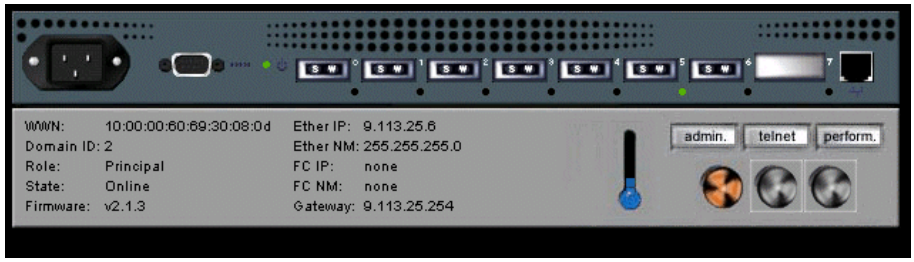

Figure 108. IBM 3534 Managed Hub

5.4.1.1 Setting the IP address using the Ethernet port

We show how to set the hub IP address by using the Ethernet port.

1. Attach the LAN to the front panel of the hub by plugging an existing Ethernet 10/100BaseT LAN cable to the RJ-45 connector on the front of the hub.
2. Turn on the hub by plugging it into an electrical outlet. Make sure that the power cord is fully seated into the front of the unit, and the green ready LED is on. Wait two minutes for diagnostics to complete.
3. From a LAN attached server, type the Telnet IP address.

   For example, issue the command: `Telnet 9.113.25.6`

If this is the initial installation, use the default IP address found on the label on the top left corner of the 3534 Managed Hub. If the 3534 Managed Hub has been installed before using the IP address on the label, continue using the current address from the label. If the IP address on the label was not used, you will need to get the current IP address from the system administrator.

After the Telnet command has been issued, the hub will respond as shown in Figure 109.

At each prompt, type in the information as shown and press Enter. The default login is Admin, and the default password is password.

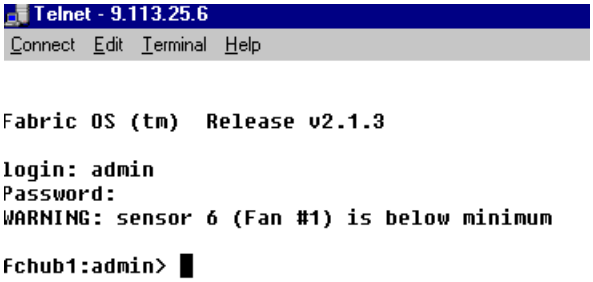

Figure 109. Start setting the IP address

4. At the prompt enter the command: `ipAddrSet`

   This is the command to set the IP address, and the result is shown in Figure 110.

   The current Ethernet IP address is shown. Now you can enter your new address. Press Enter to keep the old one.

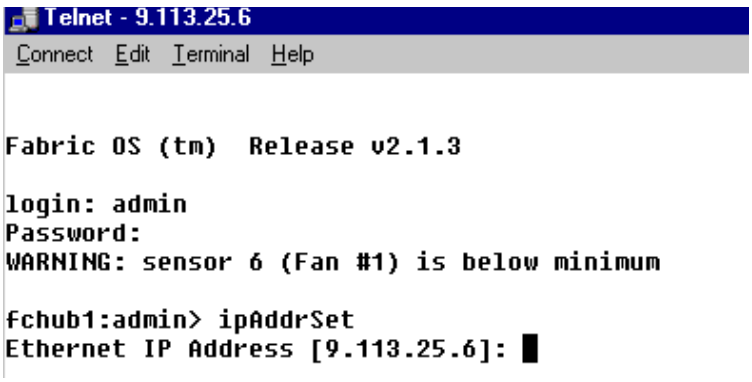

Figure 110. Current IP address

5. Ethernet Subnetmask [Current Subnetmask will be shown or None]: new Subnetmask or press Enter.

   This is the new Subnetmask from your system administrator or, if none is required, press Enter.

6. Fibre Channel IP Address [None]: press Enter.

   Fibre Channel Subnetmask [None]: press Enter.

7. Gateway Address [Current Gateway address or None]: enter new Gateway address or press Enter.

   This is the Gateway address the system administrator provided or, if none is required, press Enter.

8. Ipaddress:admin> logout

This process, in step 5 through step 8, is shown in Figure 111.

In our example we have changed the IP address from 9.113.25.6 to 9.113.25.7.

The final prompt will ask if you want to set the IP address to the new value now. Entering 'Y' installs the new value; Typing 'N' delays the change until the next hub restart.

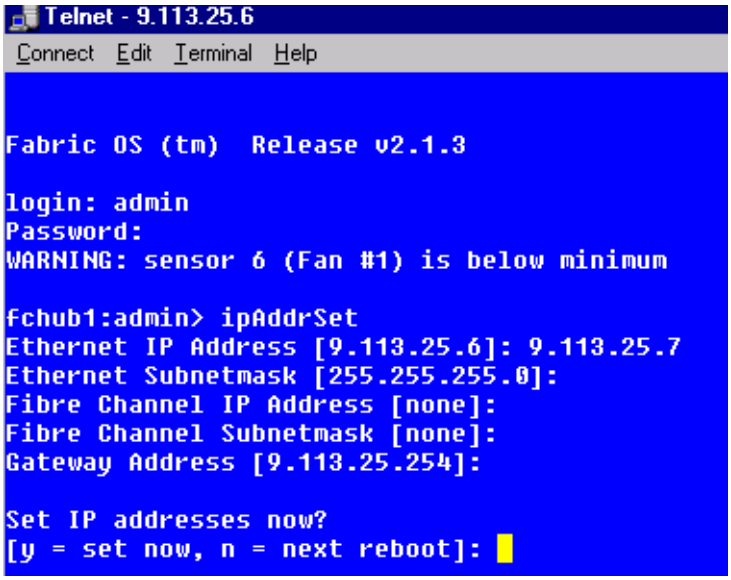

Figure 111. Changing the IP address

9. Ipaddress:admin> logout

This ends the Telnet session. We have completed the installation of the 3534 Managed Hub.

5.4.1.2 Setting the IP address using the serial port

We also show how to set the hub IP address by using the serial port.

If you also choose to set up the IP address using the serial cable, open a HyperTerminal session on your laptop, and proceed with these steps.

Note: Opening a HyperTerminal session varies depending on which version of Windows you are using. With Windows NT to start a HyperTerminal session, go to Start —> Programs —> Accessories.

Configure HyperTerminal as follows:

1. We used a laptop with Windows NT 4.0.

   Prior to plugging the hub into the electrical outlet we attached the laptop to it with the serial cable, female to female, shipped with the hub.

2. In the Connection Description window, type the name you want to use for your new session. Select any icon from the icon menu shown, and click OK.

   The Connect to window appears. In this window, change the **Connect using** setting from the default to **Direct to Com1**, and click OK.

   The COM1 Properties window appears as shown in Figure 114 on page 130.

   Set the following parameters in the Port Settings tab:

   - 8-bit
   - No parity
   - One stop bit
   - 9600 baud
   - Flow Control = None
   - Emulation = Auto Detect

3. Click File —> Properties. The Properties window appears. Select the Settings tab, set the Emulation field to **Autodetect**.

4. After this, start up the hub by inserting the power cord into electrical outlet and waiting for about two minutes for diagnostics to complete. Make sure that the power cord is fully seated into the front of the unit, and the green ready LED is on.

5. Press Enter on your laptop.

The hub responds with the prompt:

```
Admin>
```

6. The HyperTerminal session is now running. For each prompt, type in the information as shown and press Enter at the end of each response.
7. Admin> ipAddrSet

   This is the command to set the IP address.
8. Ethernet Subnetmask [Current sub net mask will be shown or None]: new Subnetmask or press Enter.

   This is the new Subnet mask from your system administrator or, if none is required, press Enter.
9. Fibre Channel IP Address [None]: press Enter.

   Fibre Channel Subnetmask [None]: press Enter.
10. Gateway Address [Current Gateway address or None]: enter new Gateway address or press Enter.

    This is the Gateway address the system administrator provided or, if none is required, press Enter.
11. Ipaddress:admin> logout

This ends the Serial port session. You have completed the installation of the 3534 Managed Hub by using the serial port. Remove the cable from the serial port connector.

5.4.1.3 Downloading Firmware

The hub is shipped with the latest level of code (firmware) available. However, new code is periodically released that you can easily download to the hub. This task requires that you save data and executable software to your server. The latest code can be obtained from the IBM SAN Fibre Channel Managed Hub, 35341RU Web site:

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

5.5 Zoning

Zoning is used to set up barriers between different operating environments to deploy logical fabric subsets by creating defined user groups, or to create test or maintenance areas, or both, which are separated within the fabric.

Zoning gives you the flexibility to manage a Storage Area Network to meet different user groups' objectives.

Zoning components are discussed in greater depth in 6.3.1, “The role of zoning in a SAN” on page 155, where the same principles apply.

The way in which to define zones, zone members, aliases and zone configurations, and the administration of zones are the same for the IBM SAN Fibre Channel Managed Hub, 35341RU, as they are for the IBM SAN Fibre Channel Switch.

5.6 Cascading

Cascading is a term used for interconnecting multiple switches.

When you start your SAN business you might begin working with a simple SAN fabric with perhaps only one switch and a few ports. But as you start implementing more and more host and storage devices in your SAN, you will reach limits where you have to add either more ports, or expand your fabric by adding more switches.

Interconnecting your first switch with other switches, allows you to build much larger fabrics. Expanding the fabric in this manner is called cascading.

The hub can be connected to one other hub to expand the loop capabilities from 7 to 14 ports. It can also be connected with a single port into a SAN fabric as a loop extension.

The attribute for interconnecting ports changes to an E_Port (expansion port).

In our hub, port 5 was used as an E_Port into an IBM SAN Fibre Channel Switch, 2109-S08. We show this in Figure 112.

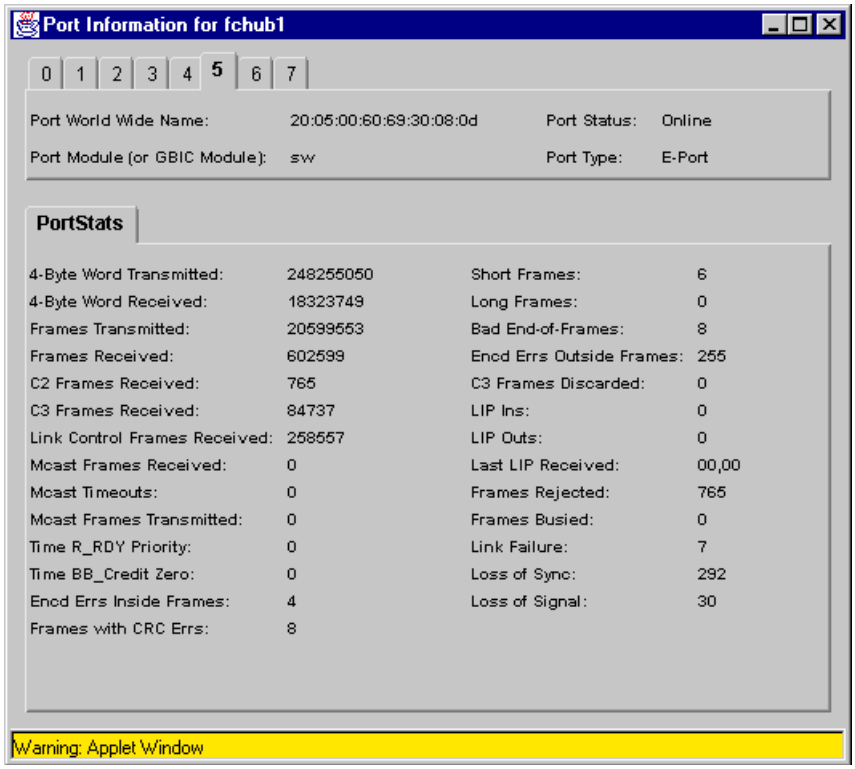

Figure 112. 3534 E_Port

5.7 QuickLoop

Fibre Channel provides three distinct interconnection topologies.The three fibre channel topologies are:

- Point-to-point
- Arbitrated loop
- Switched — referred to as a fabric

The IBM 3534 Managed Hub is a Fibre Channel Arbitrated Loop device.

A simple loop, configured using a hub, is shown in Figure 113.

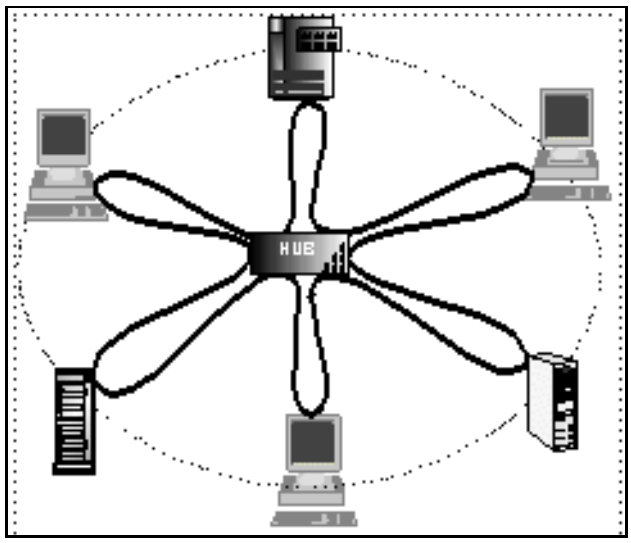

Figure 113. Arbitrated loop

There are two different kinds of loops, the private and the public loop.

5.7.1 Private loop

The private loop does not connect with a fabric, only to other private nodes using attachment points called Node Loop Ports (NL_Ports). A private loop is enclosed and known only to itself. NL_Ports can only attach to other NL_Ports or to Fabric Loop Port (FL_Ports).

In Figure 114, we show a private loop.

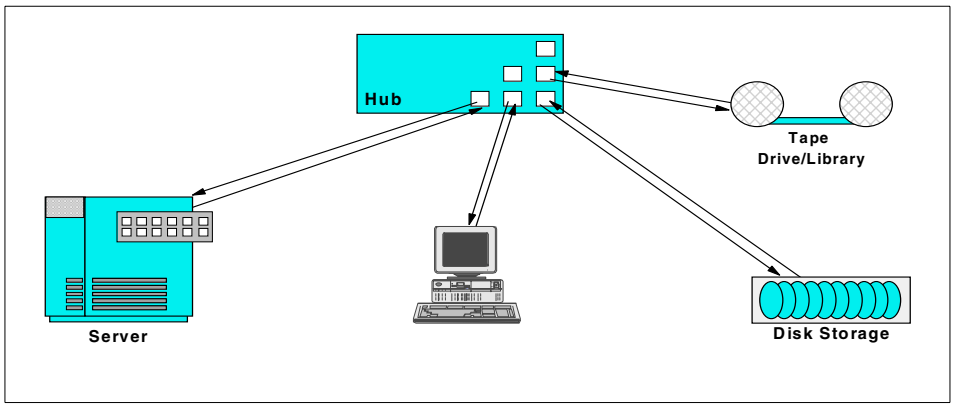

Figure 114. Private loop

5.7.2 Public loop

A public loop requires a fabric, and has at least one FL_Port connection to a fabric. A public loop extends the reach of the loop topology by attaching the loop to a fabric. FL_Ports can only attach to NL_Ports.

In Figure 115, we show a public loop.

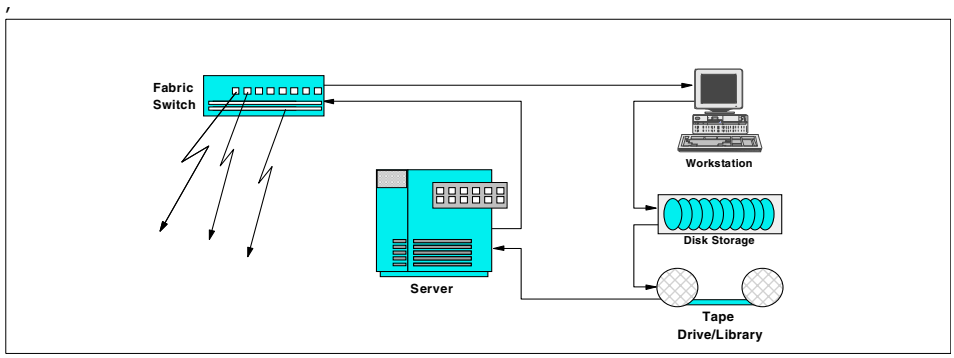

Figure 115. Public loop

5.7.3 Private devices in public fabrics

The characteristic of a fabric is allowing communication between any host or storage device in the fabric. In other words, all communication is “public”.

Problems may arise when a private Fibre Channel device is attached to a switch. Private Fibre Channel devices were designed to only work in private loops.

For more details about how node and port addressing works, refer to *Designing an IBM Storage Area Network*, SG24-5758.

5.7.4 QuickLoop feature

QuickLoop applies to both the IBM managed hub and the IBM 2109 switch. On the managed hub, QuickLoop is always in effect, on the switch, QuickLoop must be enabled at a port level.

If you have a private host (server), communication of a public nature is not possible. To solve this issue, switch vendors, including IBM, support a QuickLoop feature, or a feature which allows public to private, or private to public bridging.

QuickLoop creates a unique Fibre Channel topology, that allows host bus adapters that use Fibre Channel Arbitrated Loop (FC-AL) without knowledge

of the SAN fabric (that is to say, an adapter that cannot perform fabric login), commonly referred to as private loop initiators, to communicate with Fibre Channel Arbitrated Loop public storage devices. This can also include IBM 2109 Fibre Channel Switches.

QuickLoop allows individual switch ports to be designated as arbitrated loop ports, allowing a private loop host initiator to communicate with arbitrated loop storage devices as though they were all contained in one logical loop.

The QuickLoop feature supports legacy devices. Legacy devices are those that are in a Private Loop Direct Attach (PLDA) environment. The QuickLoop feature allows these devices to be attached to a SAN and operate no differently than in a PLDA environment.

As opposed to the IBM SAN Fibre Channel Switch, where you have the option of defining QuickLoop ports, the IBM SAN Fibre Channel Managed Hub, 35341RU ports work in QuickLoop mode as a default.

For a list of supported devices for the IBM 3534, see the Web site:

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

In a simple scenario with one host, a single hub, and some storage devices, it does not matter if we have a private or public loop or if there is a fabric.

Having a hub, rather than a switch, has an impact on the performance, and integrating the hub with QuickLoop into a SAN fabric is beneficial.

If the hub, with its arbitrated loop devices, is part of a larger SAN fabric that includes multiple 2109 switches, it is possible to expand the number of ports in the QuickLoop (called looplets) to the total of the hub ports plus the ports of one IBM SAN Fibre Channel Switch.

For example:

- 7 hub ports + 7x 2109-S08 ports = total of 14 NL_Ports
- 7 hub ports + 15x 2109-S16 ports = total of 22 NL_Ports

There are two options with regard to how a hub is employed in this manner:

- As a single-hub: all looplets of a QuickLoop reside in one hub.
- In a Hub-Switch combination: looplets of a QuickLoop span across hub plus switch.

In Figure 116, we show high-performance, multi-target connectivity to a single PLDA, where all the ports of the hub operate in QuickLoop mode. The hub serves as a concentrator.

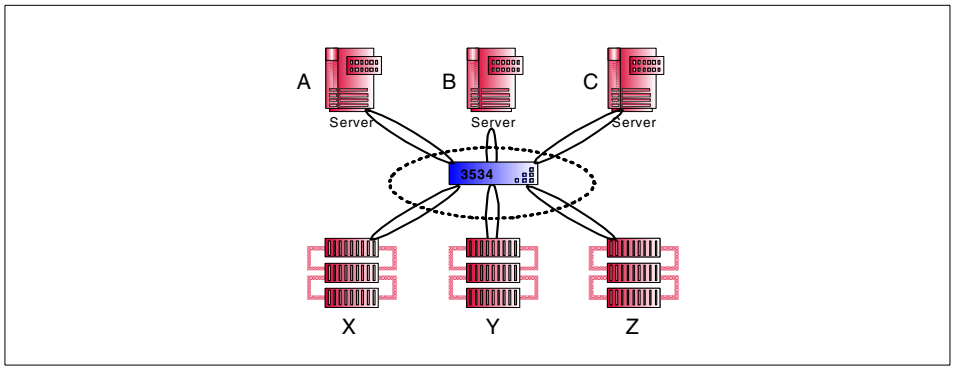

Figure 116. QuickLoop using managed hub

In Figure 117 we show how the QuickLoop feature supports legacy devices. The legacy devices refer to devices that are in a PLDA environment.

The QuickLoop feature allows these devices to be attached to a Storage Area Network and operate no differently than if they were in a PLDA environment. When QuickLoop mode is enabled, all ports defined as QuickLoop ports on the switch behave as if they are one logical Fibre Channel Arbitrated Loop.

The switch needs the QuickLoop facility available from IBM as RPQ 8S0521.

RPQ 8S0521 is a licensed product requiring a valid license key. The RPQ 8S0521 License Agreement is covered by the IBM Agreement for the Licensed Internal Code and is linked to the serial number of the hub.

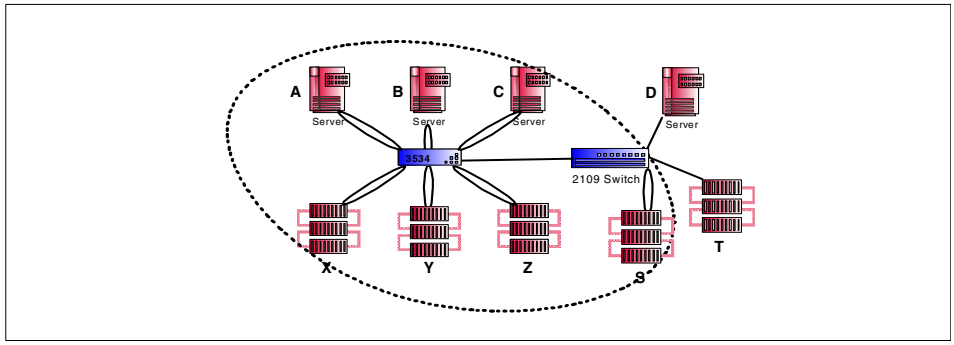

Figure 117. Quickloop spanning to switch

Both configurations allow up to 126 NL_Port devices in one QuickLoop.

Note

The IBM 3534 managed hub is an OEM product from Brocade. In Brocade terminology this is a switch.

5.7.4.1 Managing QuickLoop

To manage QuickLoop for the IBM 3534 Managed Hub you can choose from two methods:

- Using Telnet commands
- Using the managed hub StorWatch Specialist

Figure 118 shows the hub/switch combination we used in our ITSO scenario.

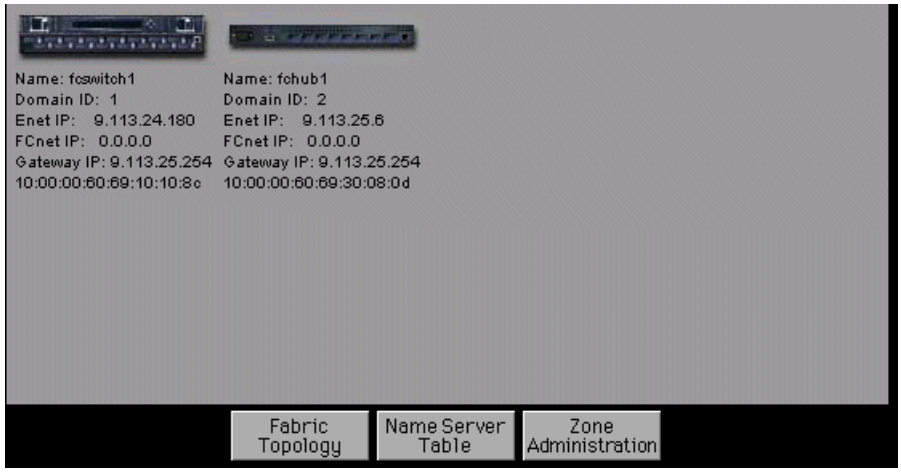

Figure 118. 3534 hub and 2109 switch

We show the Storwatch Specialist as an example of how to administer the QuickLoop. Figure 119 shows the 3534 QuickLoop entry panel.

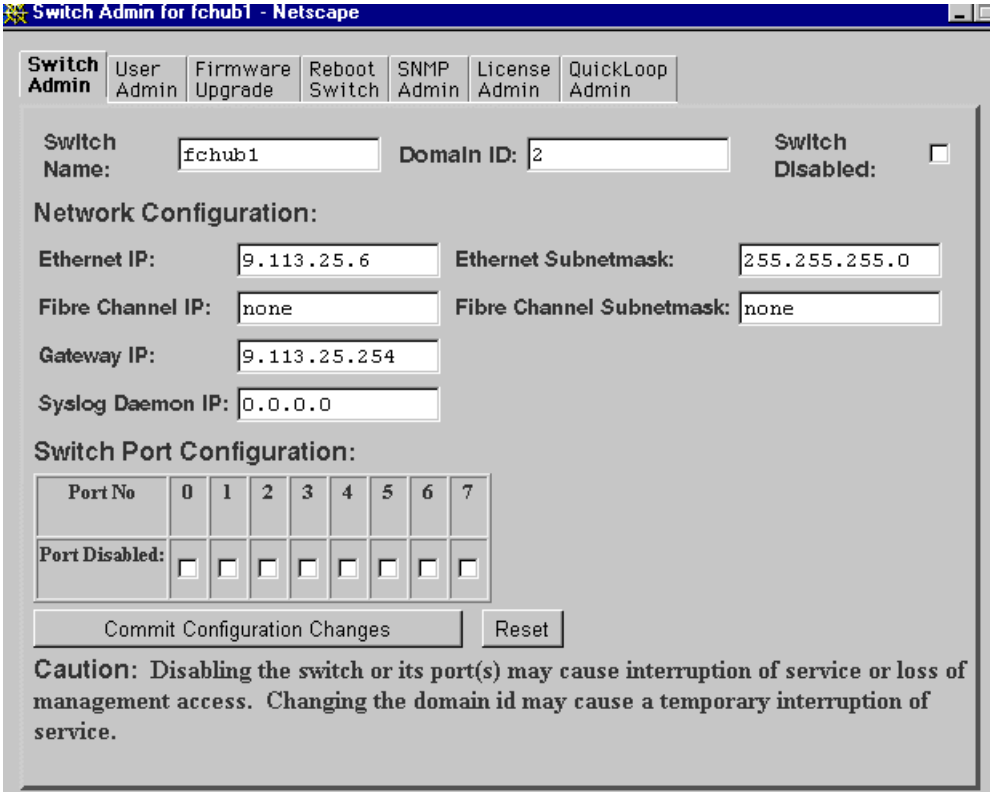

Figure 119. StorWatch QuickLoop panel

In Figure 120, we expanded the number of looplets (ports) in the private loop by bringing in a QuickLoop partner, fcswitch1.

You have two options:

- Select a QuickLoop partner and enter its WWN
- Select None

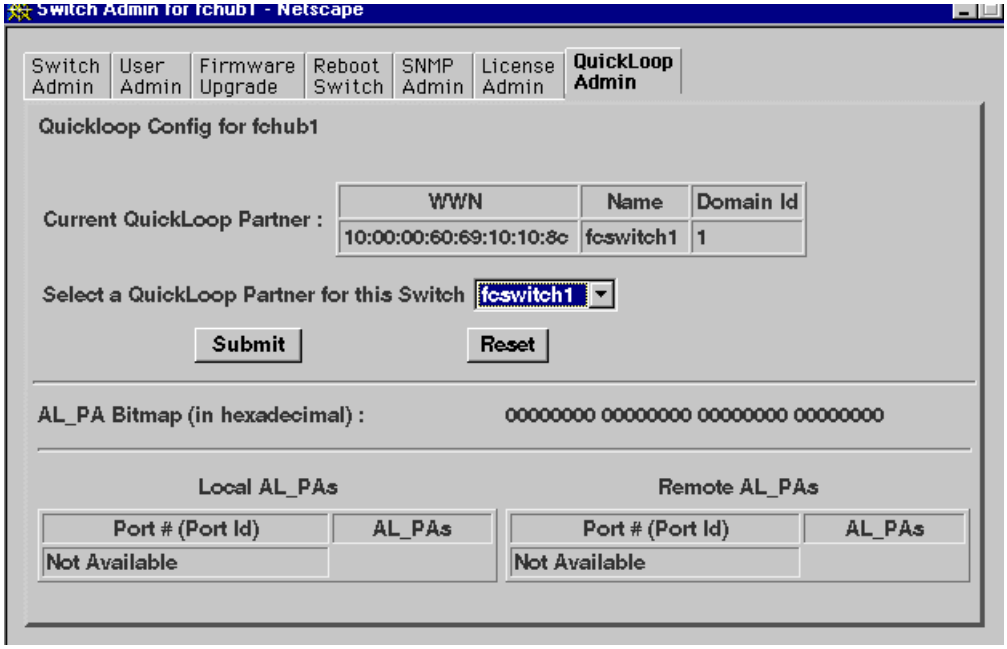

Figure 120. QuickLoop with switch partner

Telnet QuickLoop commands are also supported and these are described in *IBM 3534 SAN Fibre Channel Managed Hub User's Guide*, GC26-7391.

5.7.5 IBM 3534 managed hub zoning

The IBM SAN Fibre Channel Managed Hub, 35341RU, is able to use zoning for finer segmentation in the access of devices. Also it is used to protect devices against LIP. As the managed hub is always in QuickLoop mode, this means that all devices on that QuickLoop are subject to LIP exposure. By using zoning, it is possible to limit the effect of LIP to QuickLoop zones. That is to say, a LIP in one zone does not cause a LIP in another zone.

Zoning in the IBM 3534 works exactly as zoning does in the IBM SAN Fibre Channel Switch, and can be managed by using Telnet commands or by using the IBM Storwatch Specialist.

Zoning is described in 6.4, "Implementing zoning" on page 161.

Chapter 6. Implementing an IBM SAN Fibre Channel Switch

A SAN is a dedicated, high-speed network of directly connected storage elements, designed to move large amounts of data between host independent, distributed devices. It is a new infrastructure between servers and storage devices created to match the rapidly growing requirements of business solutions, for example, e-business and mission critical applications.

In addition to ultra high performance and 24 x 7 availability, the key requirements are improved control over storage resources and the maximum flexibility for the flow of information. The SAN storage infrastructure is clearly an enabler for this flexible information flow.

As we look at the interconnect components in a SAN installation, the Fibre Channel switch is one of its core elements. In general, they are used to implement a Fibre Channel fabric topology.

A fabric is an intelligent, shared, interconnected scheme of FC server and storage nodes. A single switch can be used to build a basic fabric. By interconnecting to other switches you can build much larger fabrics. Interconnecting switches in this way is called cascading. Cascading allows for higher availability and larger topologies.

Fibre Channel switches can be used in entry level enterprise heterogeneous implementations, and also in the largest of enterprise environments. Any Fibre Channel enabled device can be connected to any Fibre Channel switch. For flexible administration and control of storage-to-server connections, most switch vendors provide a set of management interfaces and management software.

In this chapter we show the practicalities associated with implementing an IBM SAN Fibre Channel Switch. The following topics cover:

6.1 IBM SAN Fibre Channel switch

IBM offers three different types of switches:

- The SAN Fibre Channel Switch, 2109-S08, which is an OEM product from the Brocade SilkWorm family and is an 8-port model.
- The IBM SAN Fibre Channel Switch, 2109-S16, which is an OEM product from the Brocade SilkWorm family and is a 16-port model.
- The McDATA Enterprise Fibre Channel Director, 2032-001, which is a 32-port director. Implementation of this is discussed in Chapter 7, “Implementing the McDATA ED-5000” on page 181.

In the following sections, we will show how to implement the SAN Fibre Channel Switch, 2109-S08. Figure 121 shows a picture of the 2109 model S08.

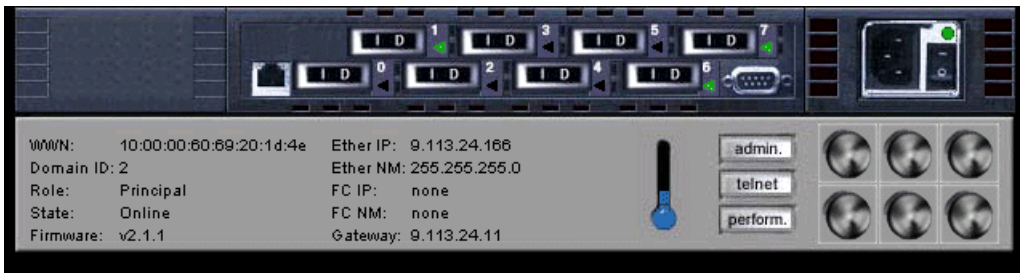

Figure 121. IBM 2109-S08 Fibre Channel switch

6.1.1 IBM 2109-S08 hardware components

The IBM SAN Fibre Channel Switch, 2109-S08, is an 8-port Fibre Channel gigabit switch that consists of a motherboard with connectors for supporting up to eight ports.

The motherboard is enclosed in an air-cooled chassis which may be a standard rack or used as a standalone unit. The chassis includes one or two power supplies, a fan tray, an RJ-45 Ethernet connection for switch set up and management, and a serial port.

Serial Port connection

The serial port is used for recovering factory settings only and for the initial configuration of the IP address for the switch, if the default address is not known. It is not used during normal operation. The IBM SAN Fibre Channel Switch, 2109-S16 does not have a serial port.

Ethernet connection

It is possible to connect an existing Ethernet 10/100BaseT LAN to the switch using the front panel RJ-45 Ethernet connector. This allows access to the switch's internal SNMP agent, and also allows remote Telnet and Web access for remote monitoring and testing. The IP address may be changed using the Ethernet Port, which is shown in Figure 122.

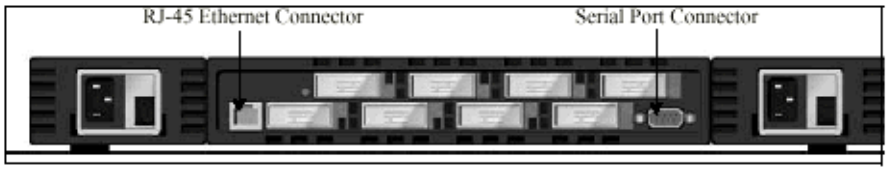

Figure 122. SAN Fibre Channel Switch, 2109-S08 Ethernet and serial connectors

In Figure 123, we show how the ports are numbered sequentially, starting with zero for the left most port. The two optional power supplies are shown to the left and right of the switch ports.

.

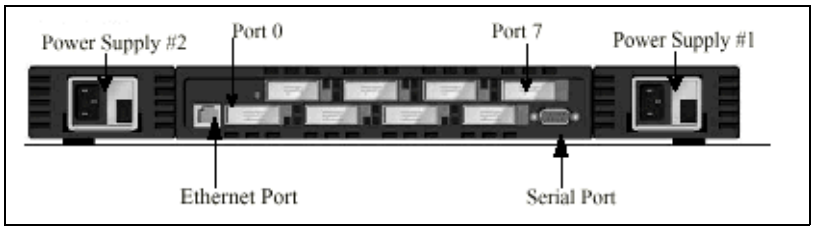

Figure 123. 2109-S08 front panel

GBICs

The switch accommodates up to eight GBIC modules. All interfaces have status lights which are visible on the front panel, giving a quick, visual check of the GBICs status and activity.

The GBIC modules supported are the Short Wave Length (SWL) and Long Wave Length (LWL) fiber-optic versions.

If you install less than eight GBIC modules, the unused port positions are protected by a metal, spring-loaded door.

Fibre Channel connections

The IBM SAN Fibre Channel Switch supports the following types of Fibre Channel connections:

- Fabric (F_Port)
- Arbitrated loop — public and private (FL_Port)
- Interswitch connection (E_Port)

Currently, only same vendor switch interconnection (E_Port) is supported.

Host attachment
The IBM SAN 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 host system we used is a Netfinity 5500 running Windows NT

Device attachment
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

We attached to the IBM Enterprise Storage Server, 2105-F20 with native FC attachment (FC 3022).

6.1.2 IBM 2109 software features

You can manage the IBM Fibre Channel Switch in three different ways:

- By using the StorWatch Fibre Channel Switch Specialist

  This is a user-friendly, Web browser interface.
- By using Telnet commands
- With SNMP

We used the Fibre Channel Switch Specialist.

The IBM Fibre Channel Switch provides advanced management capabilities for:

- Automatic discovery and registration of host and storage devices
- Intelligent rerouting of connection paths, should a port problem occur
- Cascading of switches, for scaling to larger configurations and to provide resiliency for high data availability
- Switch zoning for fabric segmentation
- Configuration with hot pluggable port GBICs for shortwave or longwave optical connections of up to 10 kilometers

Zoning is detailed in 6.3, “Zoning in an IBM SAN Fibre Channel Switch environment” on page 154, and cascading switches is detailed in 6.5, “Cascading IBM 2109 switches” on page 171.

6.2 The ITSO environment

The ITSO target SAN installation consists of:

- IBM Netfinity 5500 server running Windows NT 4.0 Service pack 5
- IBM ESS 2105-F20 with native FC adapter
- IBM 2109-S08 Switch

6.2.1 Installing the 2109-S08 switch

We recommend that you utilize a pre-installation checklist. An example of a pre-installation checklist and a description of the GBICs (and how to insert them into the switch) is detailed in the *IBM 2109 Model S08 User’s Guide*, SC26-7349 and the *IBM 2109 Model S08 Installation and Service Guide*, SC26-7350.

To install the SAN Fibre Channel Switch, 2109-S08, we followed these steps:

1. Power-On Self Test

   The switch is designed for maintenance free operation and supports Power-On Self-Test (POST) and diagnostic tests. The diagnostic tests determine the switch’s status, isolate problems, and will take approximately two minutes.

   As one of our first steps, we powered on the switch and ensured that it ran its POST successfully.

   After the switch has completed the POST, the GBIC modules return to a steady state from the flashing state that is shown during the test.

2. Setting the IP address

   The switch is shipped from the factory with a default IP address on the switch of 10.77.77.77. This IP address is noted on a label on the top front edge of the switch. This address is for the external Ethernet connection.

   In your environment, you will want to utilize your own range of IP addresses and you will have to change the default address to establish a network connection to the switch.

   Your LAN administrator should be able to provide a valid, free IP address or a range of addresses.

   We changed the default IP address to reflect the ITSO environment.

   To change an existing IP address you can either use the Ethernet port and Telnet commands entered from a server which has access to the same LAN, or you can use the switch serial port using a serial terminal emulator. The IBM SAN Fibre Channel Switch, 2109-S16, allows you to change the IP address using the control panel on the front of the switch. We document this in 6.2.3, “Setting the IP address on an IBM 2109-S16” on page 146.

In our installation we will show how to use the serial port connector shown in Figure 124, to set the IP address.

Figure 124. RJ-45 Ethernet Connector and Serial Port Connector

6.2.2 Setting the IP address using the serial port (2109-S08 only)

Here are the steps to set the IP address using the serial port:

1. Connect a system with a serial terminal emulator to the serial port (prior to powering up the switch).
2. Start up a terminal emulation session (hyperterm).
3. Power up the switch. If two supplies are present, power up both. As the switch goes through its diagnostics, it will post messages to your service terminal. When the switch is fully powered on, the switch will respond: `Admin>`.
4. `Admin> ipAddrSet` (this is the command to set the IP address); press Enter.

5. Ethernet IP address [Current ipaddress or None]: new IP addr; press Enter.

   This would be a new address provided by the Network Administrator and in our example we used 9.113.24.19.

6. Ethernet Subnetmask [Current subnet mask or None]: new Subnetmask or press Enter.

   This would be a new Subnetmask provided by the network administrator and we set ours to 255.255.255.0.

7. Fibre Channel IP Address [None]: press Enter.
8. Fibre Channel Subnetmask [None]: press Enter.
9. Gateway Address [Current Gateway address or None]: new Gateway address or press Enter.

   This would be a new Gateway address provided by the network administrator.

10. `Admin> Logout`; press Enter.

This will end the Telnet session. This completes setting the IP address.

Serial port settings

In our test scenario we used an IBM laptop as a serial terminal emulator and linked it to the switch using a standard serial cable with two female 9-pin connectors.

The serial port settings are:

- 8-bit
- No parity
- One stop bit
- 9600 baud
- Flow Control = None
- Emulation = Auto Detect

It is important to change the Ethernet Subnetmask and the Gateway address.You can see the entries on the front panel of the switch in Figure 125.

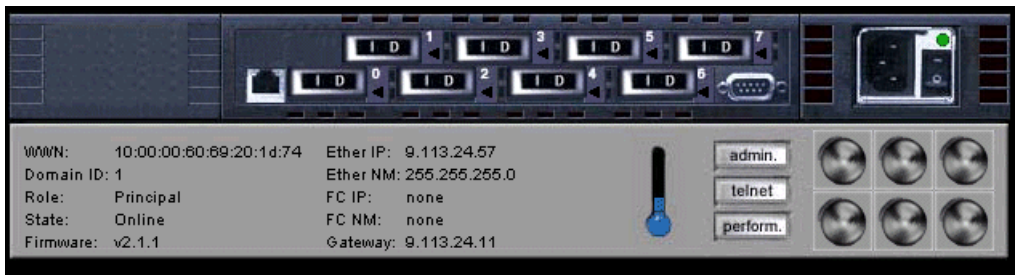

Figure 125. 2109-S08 switch

Now the switch should be linked to the network and you should be able to access it from any server in the network.

6.2.3 Setting the IP address on an IBM 2109-S16

One of the differences between the S08 and the S16, apart from the increase in port counts, is the control panel that has been introduced. There is no serial port on the S16.

From the control panel it is possible to set the IP address. The front panel of the IBM SAN Fibre Channel Switch, 2109-S16 is shown in Figure 126.

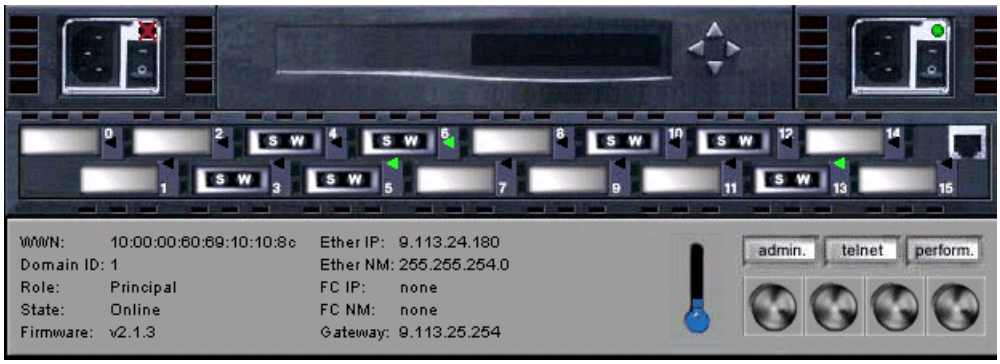

Figure 126. 2109-S16

The triangular buttons to the right of the picture, just above port 12, represent the data entry buttons referred to in the following steps:

1. Power up the switch

   Wait two minutes while the diagnostics complete. The display panel will go blank after the diagnostics are complete.

2. Press the ^ button, the switch will respond: `Configuration menu`.

3. Press the > button the switch will respond: `Ethernet IP address`.
4. Press the > button the switch will respond by:

   Displaying the current Ethernet IP address in the form: `XXX XXX XXX XXX`
5. To change the current address to the preferred address, do the following:
   - Pressing the < button, will move the cursor (the entry point) from one field to the next. If you go past a field, continuing to press the < button will cause the cursor to wrap around and return to the desired spot.
   - Pressing the ^ button or the v button will either increment or decrement the current field. Holding the button down will make this happen rapidly. The numbers in the field will wrap to 0 from 255 or from 255 to 0, depending on whether you are incrementing or decrementing. This will help you get to the desired value quickly.

   When you have all fields set to the desired value, continue with step 6.
6. Press the > button. The switch will respond: `Accept? Y N`.

   Pressing the < button indicates Yes. Pressing the > button indicates No.
   - If you respond no, the switch will again respond: `Ethernet IP address`. You can now restart the process.
   - If you respond yes, the switch will respond: `updating the Config`.

   This will cause the new address to be stored. After the switch has made the change, it will again display: `Ethernet address`.

   If no other address is to be changed (Gateway, Subnet, etc.), you may simply stop here by pressing the < button. At this point, you are done setting up the switch. If you need to set other switch addresses (Gateway, Subnetmask), press the ^ button and the switch will respond: `Ethernet Subnetmask`.
7. Press the > button, the switch will respond by displaying the current Ethernet Subnetmask in the form: `XXX XXX XXX XXX`.
8. To change the current Subnetmask to the preferred address do the following:

   Pressing the ^ button or the v button will either increment or decrement the current field. Holding the button down will make this happen rapidly. The numbers in the field will wrap to 0 from 255 or from 255 to 0 depending on if you are incrementing or decrementing. This can help you get to the desired value quickly. When you have all fields set to the desired value, continue with step 9.
9. Press the > button. The switch will respond: `Accept? Y N`.

Pressing the < button will indicate Yes. Pressing the > button will indicate No.

- If you say No, the switch will again respond: `Ethernet Subnetmask`, and you can restart the process.
- If you say yes, the switch will respond: `Updating the Config`.

This will cause the new address to be stored. After the switch has made the change, it will again display: `Ethernet Subnetmask`.

- If no other address is to be changed (Gateway, Subnet, etc.), you may simply stop here by pressing the < button. At this point, you are done setting up the switch.
- If you need to set other switch addresses (Gateway, Domain), press the ^ button and the switch will respond: `Fibre Channel IP address` (will not be required at this time).

10. Press the ^ button and the switch will respond: `Fibre Channel Subnetmask` (will not be required at this time).
11. Press the ^ button and the switch will respond: `Gateway Address`
12. Press the > button and the switch will respond by displaying the current Gateway address in the form: `XXX XXX XXX XXX`.
13. To change the current Gateway address to the preferred address:

    - Press the < button to move the cursor (the entry point) from one field to the next.

    If you go past a field, continue to press the < button to wrap the cursor around and return to the desired spot.

    - Pressing the ^ button or the v button, will either increment or decrement the current field. Holding the button down will make this happen rapidly. The numbers in the field will wrap to 0 from 255 or from 255 to 0 depending on if you are incrementing or decrementing. This can help you get to the desired value quickly.

    When you have all fields set to the desired value, continue with step 14.
14. Press the > button and the switch will respond: `Accept? Y N`.

    Pressing the < button will indicate Yes. Pressing the > button will indicate No.

    - If you say no, the switch will again respond: `Gateway Address`, and you can restart the process.
    - If you say yes, the switch will respond: `Updating the Config`.

This will cause the new address to be stored. After the switch has made the change, it will again display: `Gateway Address`.

- If no other address is to be changed (Gateway, Subnet, etc.), you may simply stop here by pressing the < button. At this point you are done setting up the switch.
- If you need to set other switch addresses (Domain), press the ^ button. The switch will respond: `Domain`.

15. Press the > button the switch will respond by: Displaying the current Domain in the form: `XXX XXX XXX XXX`.

    To change the current Domain to the preferred address:

    - Press the ^ button or the v button to increment or decrement the current field. Holding the button down will make this happen rapidly. The numbers in the field will wrap to 0 from 255 or from 255 to 0 depending on whether you are incrementing or decrementing. This can help you get to the desired value quickly.

    When you have all fields set to the desired value, continue with step 16.

16. Press the > button and the switch will respond: `Accept? Y N`.

    Pressing the < button will indicate Yes. Pressing the > button will indicate No.

    - If you say no, the switch will again respond: `Domain`, and you can restart the process.
    - If you say yes, the switch will respond: `Updating the Config`.

    This will cause the new address to be stored.After the switch has made the change, it will again display: `Domain`.

17. Respond by pressing the < button.

You have completed installation of the 2109 S16 switch. To perform a quick check of the switch's Fibre Channel ports, follow the switch installation verification process.

Switch installation verification

1. Power down the switch.
2. Power up the switch (if dual power, power up both supplies).
3. Verify that the associated power supply LEDs are on.
4. Wait two minutes while POST diagnostics run.
5. Verify that the switch ready LED is on.

6. Plug the appropriate wrap connector (black for shortwave and gray for longwave) into each GBIC. Verify that each associated port LED shows a slow green (every two seconds) flash.

If any of the above checks have failed, see Problem Determination Start Map on page. If none fail, the switch is ready for use.

6.2.4 Attaching the switch to a network server and a storage device

Now that the switch is working, attach one of the eight available ports, using a Fibre Channel cable, to your host server Fibre Channel adapter; and another switch port, again using a Fibre Channel cable, to your storage device.

Depending on the GBIC in the port, this can be either an SWL fiber-optic link or an LWL fiber-optic link.

The SWL fiber-optic GBIC module, with SC connector color-coded black, is based on shortwave length lasers supporting 1.0625 Gb/s link speeds. This GBIC module supports 50-micron, multi-mode, fiber-optic cables, with cables up to 500 meters in length.

The LWL fiber-optic GBIC module, with SC connector color-coded blue, is based on longwave length 1300nm lasers supporting 1.0625 Gb/s link speeds. This GBIC module supports 9-micron, single-mode fiber. Cables up to 10 kilometers in length with a maximum of five splices can be used.

6.2.4.1 Attaching the server and the storage device

To establish a SAN with fabric point-to-point connection between server and storage, your server needs a Fibre Channel host bus adapter (HBA) card that supports fabric point-to-point.

In our environment, we used the following:

- Netfinity 5500NT with an EMULEX LP7000 FC card
- Netfinity and the ESS with standard FC cables connected to the switch
- StorWatch Enterprise Storage Server Specialist to add our Netfinity host to the other hosts already defined to the ESS

Support matrix

A host and storage support matrix for the IBM 2109-S08/S16 switches with host bus adapter information can be found on the Web at:

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

The latest list of ESS supported servers also gives you host adapter card information. This is found on the Web at:

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

6.2.4.2 Defining a host to the ESS

This is described in 4.1.4, “Defining a new host FC port with its WWPN” on page 80, and the following topics assume that this has been performed.

The nickname that we used for the host is *ITSO*.

6.2.4.3 Configuring the host adapter

This has been described in 4.1.7, “Configuring an ESS Fibre Channel port” on page 96, and we will assume that this has been performed and represents the topology that is employed in your environment.

Remember that for a SAN fabric solution, you must select Fibre Channel Point to Point on the ESS.

6.2.4.4 Assigning array capacity

This is described in 4.1.5, “Configuring disk groups” on page 84.

6.2.5 SAN installation verification

If everything has been configured correctly, and all the parameters are set appropriately, this entry-level SAN installation should now run as a Fibre Channel Point to Point, fabric solution.

We used the StorWatch Fibre Channel Switch Specialist to verify our SAN solution.

6.2.5.1 Launching the StorWatch Fibre Channel Switch Specialist

Access to the IBM StorWatch Specialist is provided through one of the following Java-enabled Web browsers.

- For Windows 95/98 or Windows NT:
  - Internet Explorer 4.0 or above
  - Netscape 4.51 or above
- For UNIX:
  - Netscape 4.51 or above

In addition to the above, Java Plug-In 1.3.0 is recommended.

To Launch

1. Start the Web browser, if it is not already active.

2. Enter a switch name or IP address in the Location/Address field.
3. The Fabric View appears, displaying all compatible switches in the fabric.

We show an example of a cascaded fabric in Figure 127.

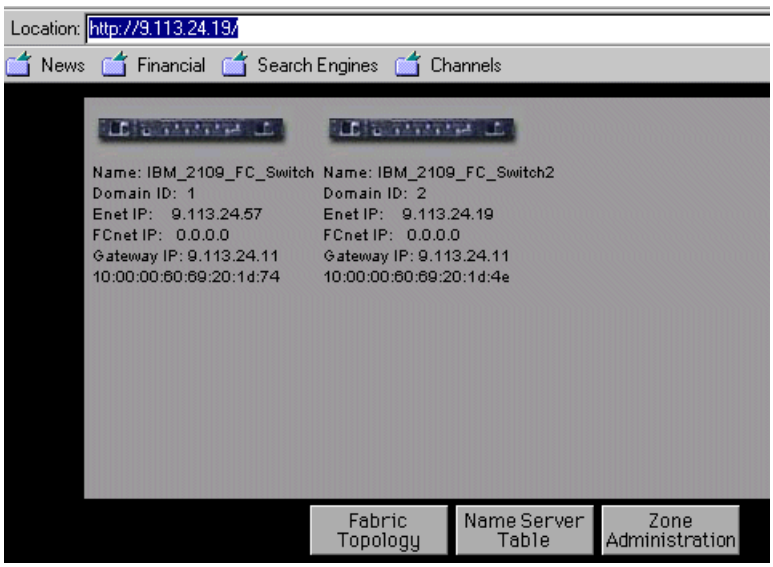

Figure 127. Launch view

In our installation, port 1 of the switch was connected to the Netfinity NT host and port 6 was connected to the ESS storage device. The LED to the right of ports 1 and 6 is a steady green (indicating link is working) and shown in Figure 128.

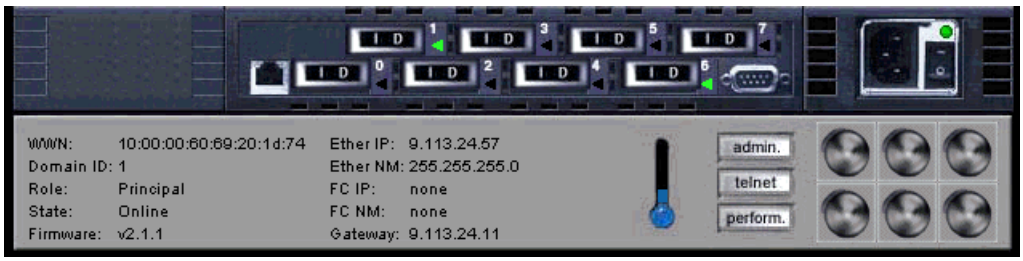

Figure 128. Port attachment

In Figure 129 we indicate the F-Port connection of port 6, which is attached to the storage device. As our Port Type is ‘F-Port’, indicating fabric, we know that our installation has been successful.

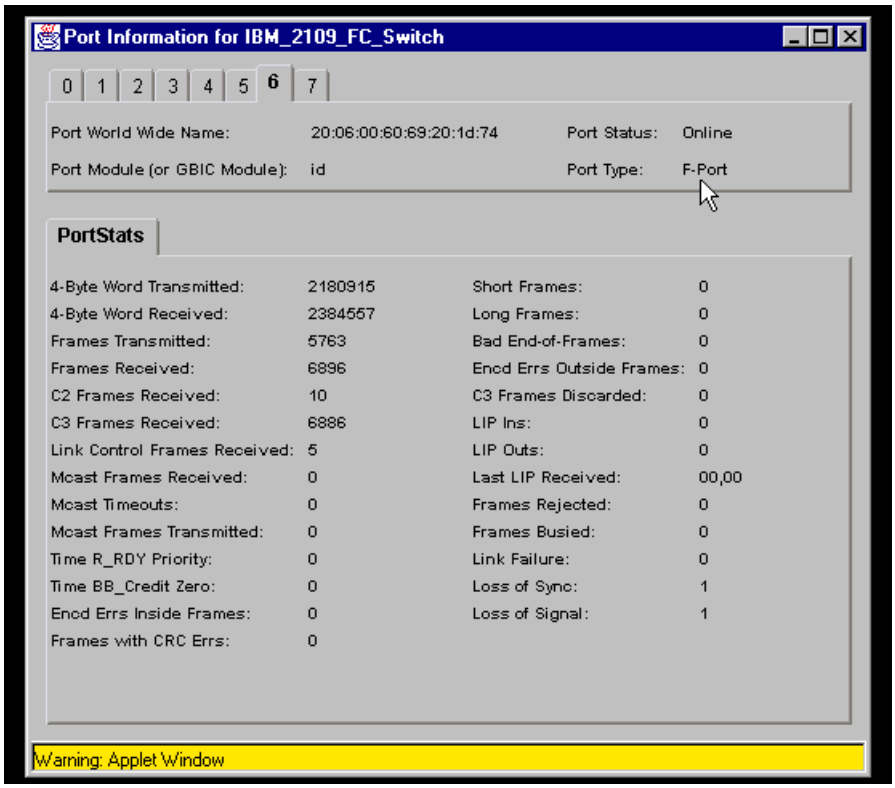

Figure 129. F-Port storage connection

In Figure 130, we show the operation mode of port 1, which is the host connection, as L-Port. This indicates arbitrated loop and is not what we expect to see.

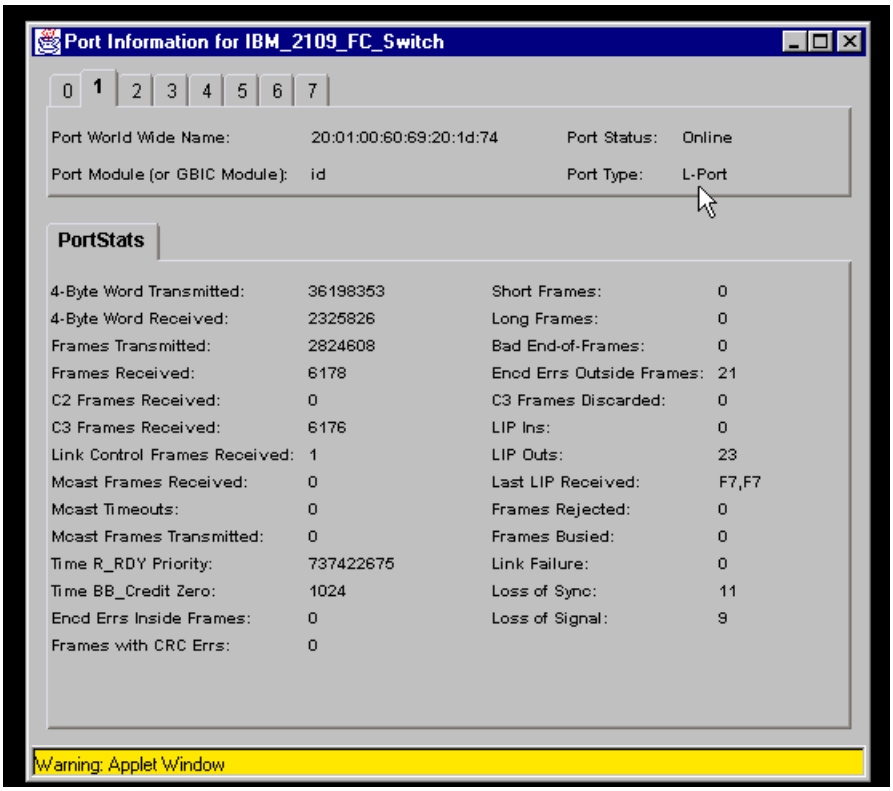

Figure 130. L-Port host connection

Our Netfinity host is working in arbitrated loop mode and we need to modify this to reflect our environment. Depending on the host fibre channel adapter cards used, a similar operation to this may have to be performed. The vendor will have to be consulted to select the correct steps for your installed adapter card.

In our case we were using the EMULEX LP7000, so we referred to that vendors documentation in conjunction with our switch installation instructions, to configure our adapter to support F-Port.

6.3 Zoning in an IBM SAN Fibre Channel Switch environment

At this stage of the chapter, we have implemented a simple SAN consisting of one host, one switch, and an ESS, connected through point-to-point Fibre Channel.

Now, we use Zoning to allow for finer segmentation of the access to storage devices from different hosts.

In these sections, we describe:

6.3.1 The role of zoning in a SAN

Zoning allows for finer segmentation of a switched fabric. Zoning can be used to establish a barrier between different environments. Only the members of the same zone can communicate within that zone and all other attempts from outside are rejected.

In our SAN scenario the idea was to limit storage access from one of the servers to only part of the shared storage facility (ESS).

Zoning may be used in your environment to create:

- Closed user groups
- Simplify resource utilization
- Facilitate time sensitive functions
- Secure fabric areas

It is obvious that a growing SAN fabric with increasing complexity needs a segmentation and control tool like zoning.

For example, it may be desirable to separate a Windows NT/2000 environment from a UNIX environment. This is very useful because of the manner in which Windows attempts to claim all available storage for itself. Because not all storage devices are capable of protecting their resources from any host seeking its available resources, it makes sound business sense to protect the environment in another manner.

By looking at zoning in this way, it can also be considered as a security feature and not just for separating environments. Zoning can also be used for test and maintenance purposes. For example, not many enterprises will mix their test and maintenance environments with their production environment. Within a fabric, you can easily separate your test environment from your production bandwidth allocation on the same fabric using zoning.

We show an example of zoning in Figure 131.

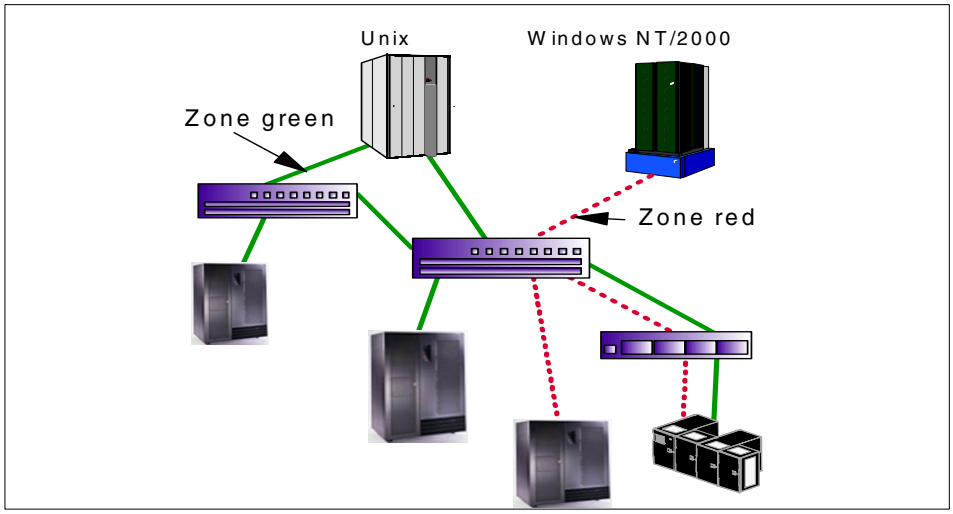

Figure 131. Zoning example

Zoning is a fabric management service used to create logical device subsets within a Storage Area Network, and enables resource partitioning for management and access control.

A zone can also be defined as a set of devices that access one another. All devices connected to a fabric may be configured into one or more zones. Devices in the same zone can see each other; devices that are in different zones cannot.

Before you can start using zoning, and before we start setting up zones in our SAN scenario, we have to explain some of the basic naming conventions and components necessary when working with zones in an IBM SAN Fibre Channel Switch environment using the StorWatch Fibre Channel Switch Specialist.

6.3.2 Zoning components

Zoning has several components besides the zones themselves, such as:

- Zone members
- Zone aliases
- Zone configurations

These components are generically referred to as zone objects.

Every zone has a name that begins with a letter and is followed by any number of letters, digits, and the underscore character “_”. Names are case sensitive, for example, Zone_1 and zone_1 are different zones. Note that spaces are not allowed.

Every zone has a member list, consisting of one or more members (empty zones are not allowed).

The maximum number of zones and the maximum number of members in a zone are constrained by memory usage. Since these limits are far larger than the number of devices connected to a fabric, they are effectively unlimited.

Zone definitions are persistent. That is the definition remains in effect across reboots and power cycles until it is deleted or changed.

A device may be a member of multiple zones.

6.3.3 Methods of implementing zones

Zoning can be implemented in two ways:

- Hardware zoning
- Software zoning

6.3.3.1 Hardware zoning

Hardware zoning is based on the physical fabric port number. The members of a zone are physical ports on the fabric switch. It can be implemented in the following configurations:

- One to one
- One to many
- Many to many

A single port can also belong to multiple zones. We show an example of hardware zoning in Figure 132.

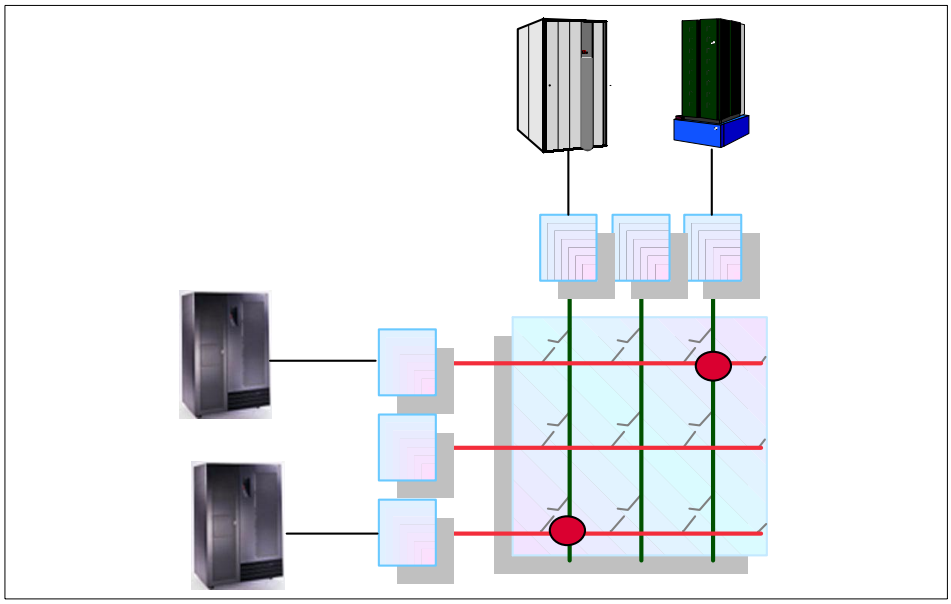

Figure 132. Hardware zoning

A physical fabric port number notation is specified as a pair of decimal numbers, where:

s - is the switch number (domain ID)
p - is the switch port number

For example, *2,12* specifies port 12 on switch number 2. When a zone member is specified by a physical fabric port number, then any and all devices connected to that port are in the zone. If this port is an arbitrated loop, then all loop devices are in the zone.

One of the disadvantages of hardware zoning is that devices have to be connected to a specific port, and the whole zoning configuration can become unusable when the device is connected to a different port. In cases where the device connections are not permanent, the use of software zoning is recommended.

The advantage of hardware zoning is that it can be implemented into a routing engine by filtering. As a result, this kind of zoning has a very low impact on the performance of the routing process.

6.3.3.2 Software zoning

Software zoning is implemented within the Simple Name Server (SNS) running inside the fabric switch. When using software zoning, the members of the zone can be defined with:

- Node WWN
- Port WWN

6.3.3.3 Zone members

Usually zoning software also allows you to create symbolic names for the zone members and for the zones themselves.

A World Wide Name notation (node and port) is specified as an eight byte hexadecimal number separated by colons. An example of this is:

```
10:00:00:60:69:00:00:8a
```

Zoning has no *field* knowledge within a World Wide Name, the eight bytes are simply compared with the node and port names presented by a device in a login frame (FLOGI or PLOGI). When a zone member is specified by node name, then all ports on that device are in the zone. When a zone member is specified by port name then only that single device port is in the zone.

The type of zone members used to define a zone may be mixed and matched. For example, a zone defined with the following members:

2,12; 2,14; 10:00:00:60:69:00:00:8a

This will contain whichever devices are connected to switch 2, ports 12 and 14, and the device with either node name or port name of 10:00:00:60:69:00:00:8a to whichever port in the fabric it is connected.

The number of members possible in a zone is limited only by the amount of memory in the fabric switch. A member can belong to multiple zones. You can define multiple sets of zones for the fabric, but only one set can be active at any time. You can activate another zone set any time you want, without the need to power down the switch.

With software zoning there is no need to worry about the physical connections to the switch. If you use WWNs for the zone members, even when a device is connected to another physical port, it will still remain in the same zoning definition because the device's WWN remains the same.

There is a potential security leak with software zoning. When a specific host logs into the fabric and asks for available storage devices, the SNS will look into the software zoning table to see which storage devices are allowable for

that host. The host will only see the storage devices defined in the software zoning table. But, the host can also make a direct connection to the storage device, while doing device discovery, without asking SNS for the information it has.

6.3.3.4 Zone aliases

Zone aliases simplify repetitive port numbers, entries, or World Wide Names. A zone alias is a C-style name for one or more port numbers or World Wide Names. For example, the name *Production* can be used as an alias for 10:00:00:60:69:00:00:8a.

6.3.3.5 Zone configurations

A zone configuration is a set of zones. At any one time, zoning may be disabled or one zone configuration may be in effect. When a zone configuration is in effect, all zones that are members of that configuration are in effect. You select which zone configuration is currently in effect.

The set of zone configurations defined in a fabric, may not be the same as the configuration that is currently in effect, and may also not be the same as the configurations that are saved in the switches flash memory. The following three terms are used to differentiate between these configurations:

- Defined configuration

  The defined configuration is the complete set of all zone objects that have been defined in the fabric. There may be multiple zone configurations defined (although only one can be in effect at a time). There may be inconsistencies in the definitions, there may be zones or aliases that are referenced but are not defined, there may be duplicate members. The defined configuration is the current state of the administrator's input.

- Effective configuration

  The effective configuration is a single zone configuration that is currently in effect. The devices that an initiator sees are based on this configuration.

  The effective configuration is built when you enable a specified zone configuration. This configuration is *compiled* by checking for undefined zone names, or zone alias names, or other inconsistencies by expanding zone aliases, removing duplicate entries, and then building the effective configuration.

- Saved configuration

  The saved configuration is a copy of the defined configuration plus the name of the effective configuration which is saved in flash memory by the *cfgSave* command. There may be differences between the saved

configuration and the defined configuration, if you have modified any zone definitions and have not saved these.

The saved configuration is automatically reloaded by the switch during power up, and if a configuration was in effect when it was saved, the same configuration is reinstated with an automatic *cfgEnable* command.

6.4 Implementing zoning

In the following sections, we cover these topics:

- 6.4.1, “Adding a second host to the ESS” on page 161
- 6.4.2, “Setting up zones in your SAN fabric” on page 162

6.4.1 Adding a second host to the ESS

Once the 2109-S08 switch is correctly implemented, adding a second host to the ESS is straight-forward from the point of view of the switch.

To connect the new host to the switch you need at least one free switch port, with the appropriate GBIC module (SWL or LWL), and a multi-mode 50/125 or 62.5/125 type cable allowing distances up to 500 meters or a single-mode 9/125 type cable with distances up to 10 kilometers.

You must use the StorWatch Enterprise Storage Server Specialist to add a new host to the ESS. This has been described comprehensively in 4.1.4, “Defining a new host FC port with its WWPN” on page 80, and more detailed information about how to do this is found in the *ESS Web Interface User’s Guide*, SC26-7346.

6.4.1.1 Adding a second host — quick steps

1. From the Storage Allocation panel, click Open Systems Storage.
2. From the Open Systems Storage panel, select Modify Host Systems.
3. To add a host system, enter the following information in the Host Attributes table, on the left side of the panel:
   a. Enter the host name in the Nickname field. A host nickname should not exceed 29 characters.
   b. Enter the host type in the Host Type field.

      Use the drop-down list box to select a host type.
4. Enter the host attachment in the Host Attachment field. Use the drop-down list box to select either SCSI attached or Fibre Channel attached. SCSI is the default.

5. Enter the World Wide Node Name in the World Wide Node Name field.

   Enter exactly 16 hexadecimal characters (0–9, a-f, and A-F) in the WWNN field. This field is mandatory for Fibre Channel hosts; it is enabled only when you select Fibre Channel in the Host Attachment field.

6. Enter a description of the host in the Description field.

   This is an optional field. The description should not exceed 255 characters.

7. Enter the remote ESS in the Remote ESS field in Import/Export Options, on the left side of the panel.

8. Click Add to add the host to the Host Systems List, on the right side of the panel.

   Note: The second host we brought to our basic SAN installation was an RS/6000 Fibre Channel AIX.

6.4.2 Setting up zones in your SAN fabric

As we stated earlier, there are several ways that zoning and zone management can be performed:

- By using Telnet using either out-of-band or in-band communication by logging into the IBM SAN Fibre Channel Switch
- By using the StorWatch Fibre Channel Switch Specialist

6.4.2.1 Using the StorWatch Fibre Channel Switch Specialist

We show how to implement zoning using the StorWatch Fibre Channel Switch Specialist. We used the following steps:

1. We started a Java-enabled Web browser and entered a switch name or IP address in the Location/Address field. We got a fabric view similar to that shown in Figure 133.

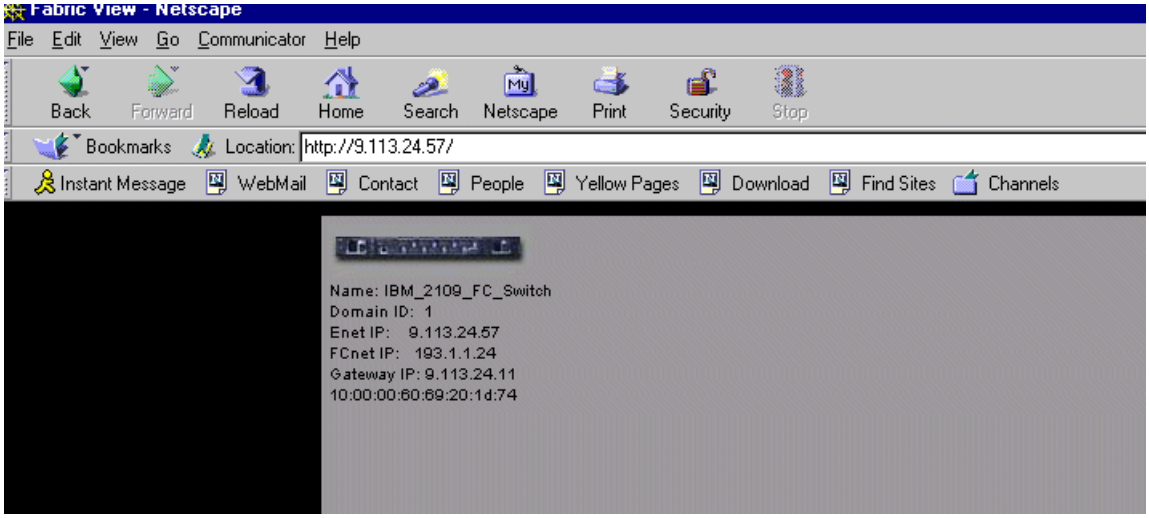

Figure 133. Fabric view

2. From this view, we clicked the Zone administration button at the bottom. Once this was done, we were presented with the view shown in Figure 134.

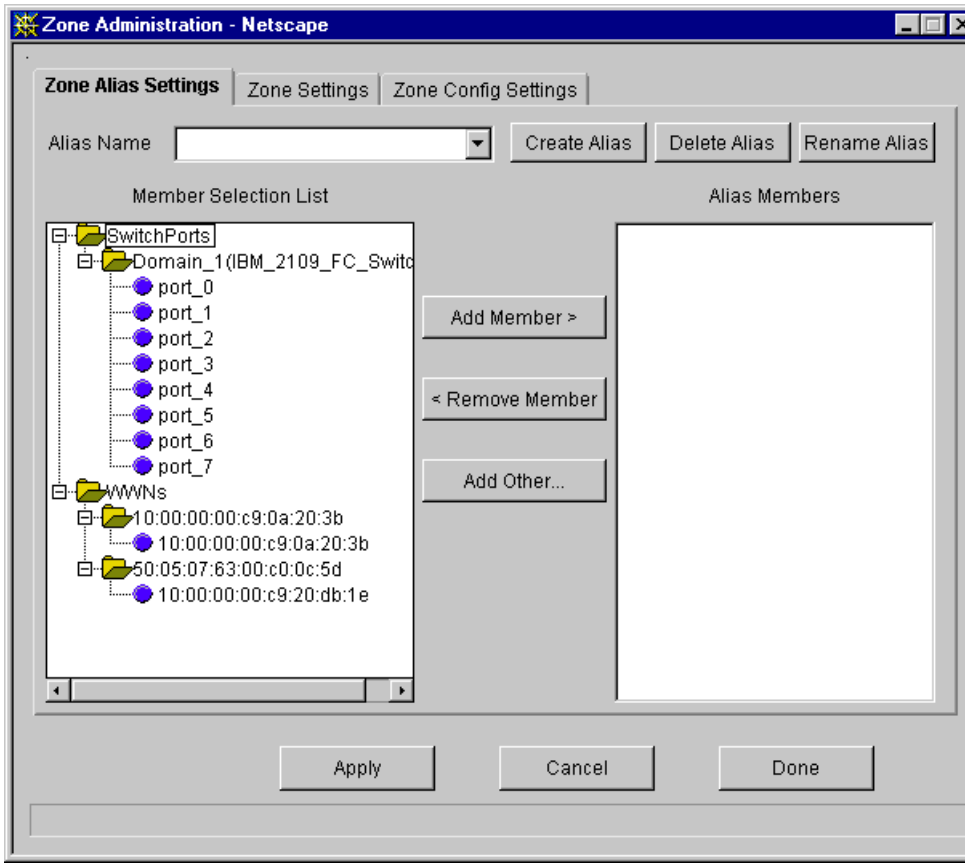

Figure 134. Zone alias settings

From this panel we can create zones in a direct way, using Zone Settings or Zone Alias Settings, to create aliases. Zone aliases are optional.

When familiarizing yourself with zoning, it may be easier to start with Zone Settings, instead of Zone Alias Settings.

3. We selected Zone Settings by clicking the middle tab. We were presented with the view shown in Figure 135.

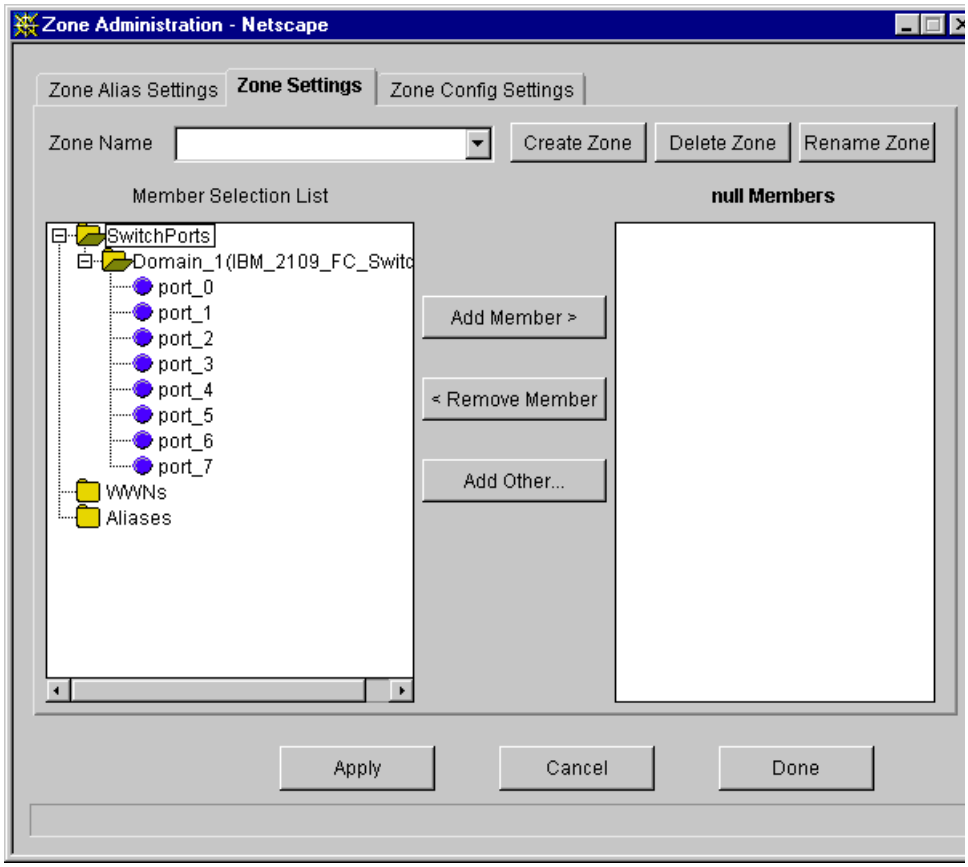

Figure 135. Zone settings view

As you can see, there are no names in the Zone Name field, and we have no entries in the member field, in the right hand panel. The 8-ports of the switch are shown, and the WWN field will contain at least one entry which is the WWN of the switch itself.

4. We created the first zone and named it “NTzone”, restricting storage access to the Netfinity host. This is shown in Figure 136.

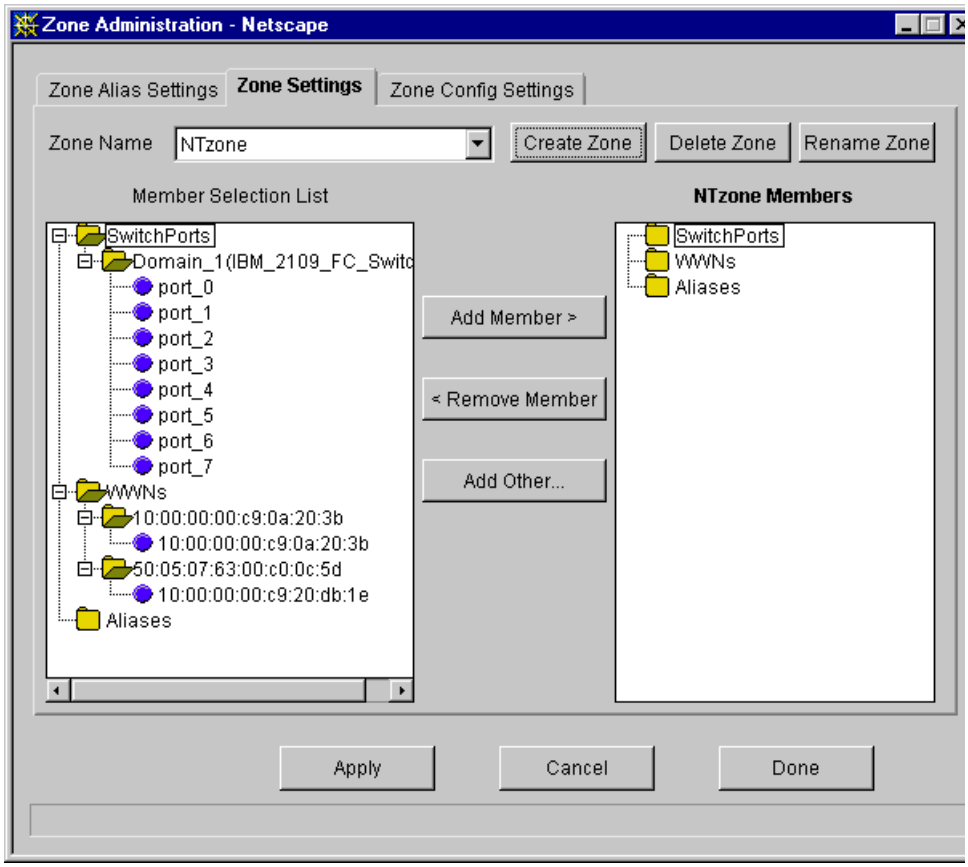

Figure 136. NTzone creation

5. As soon as we defined a zone, the system asked for zone members, because an empty zone is not allowed.

   We decided to opt for hardware zoning and allocated switch ports 0 and 1 to the NTzone. However, these members can be switch ports, WWNs, or aliases.

   We show our hardware zoning example in Figure 137. You will also notice that in the member selection list that there has been a new domain added called Domain_2. This is because the dynamic environment which we were working in had, at the time of screen capture, another switch connected to our switch (Domain_1). This is a good opportunity for us to point out that if a switch is cascaded then we have the ability to zone that into our environment too.

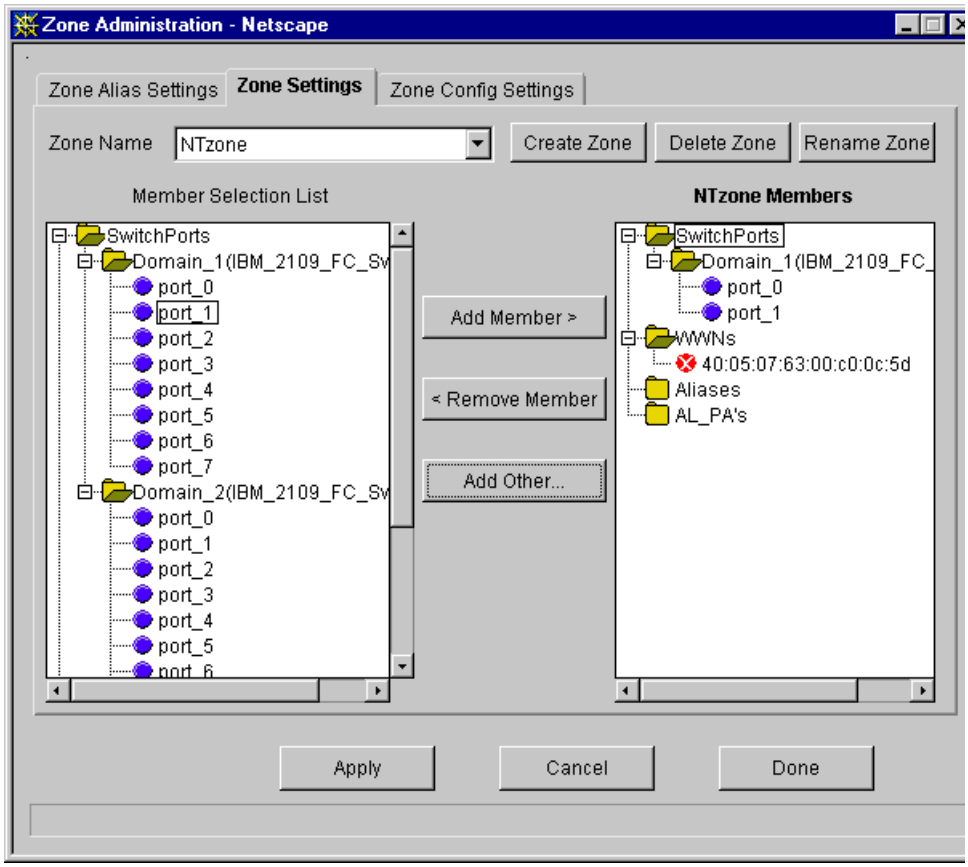

Figure 137. Choosing zone members

The NTzone contains the devices which are connected to ports 0 and 1 of the switch.

Additionally, to show how these can be used in combination with each other, we also selected the WWN 40:05:07:63:00:c0:0c:5d as a member in our NTzone.

Now, the NTzone also contains the device with the WWN (node name or port name) 40:05:07:63:00:c0:0c:5d, to whichever port in the fabric it is connected.

The WWN 40:05:07:63:00:c0:0c:5d was actually the node WWN of a tape storage device, which was not online at the time of its addition.

We have successfully defined a zone called NTzone.

6.4.2.2 Adding additional zones

To illustrate the zoning concept, we will establish a second zone and name it AIXzone. Again, we will use hardware zoning and we will use ports 2 and 3.

As a further illustration we will also add, for the Sun host, a zone named SUNzone, and we will use ports 4 and 5.

Ports 6 and 7 have been left free, because we want them for future use as E_Ports when we add (cascade) a second switch. This is covered in 6.5.1, “Cascading overview” on page 172.

Using the same process that we followed to add our initial NTzone, now we have created:

- The Netfinity zone = NTzone (ports 0,1 and WWN)
- The RS/6000 zone = AIXzone (ports 2,3)
- The SUN zone = SUNzone (ports 4,5)

Our zone names are shown in the Zone Name field, and in Figure 138, we show the SUNzone members.

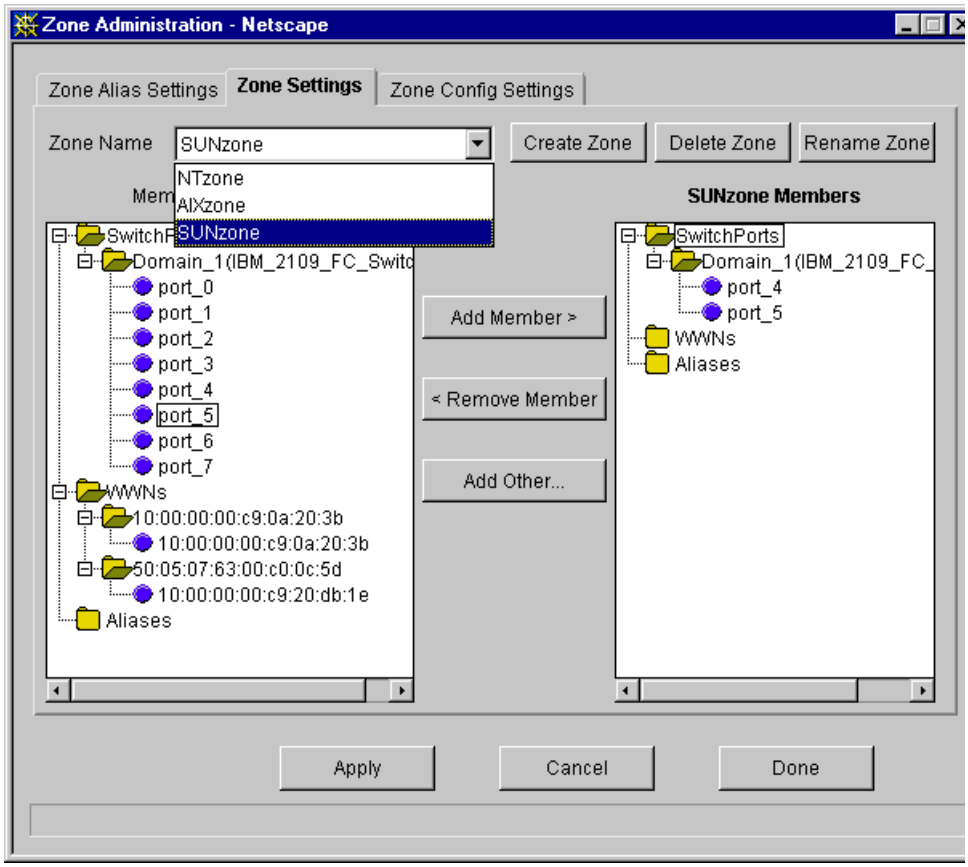

Figure 138. Creating zone names

6.4.2.3 Zone aliases

Zone aliases simplify repetitive port numbers, entries, or World Wide Names. A zone alias is a C-style name for one or more port numbers or World Wide Names. For example, the name “ESS” can be used as an alias for 10:00:00:00:c9:20:db:0c

From the panel shown in Figure 139, click the Zone Alias Settings tab, enter an Alias Name, and click the Create Alias button to create an alias. Use the Add Member field to add, or use the Remove Member field to remove an alias.

In Figure 139, we show the alias window for ESS, which we have defined as an alias for WWN 10:00:00:00:c9:20:db:0c.

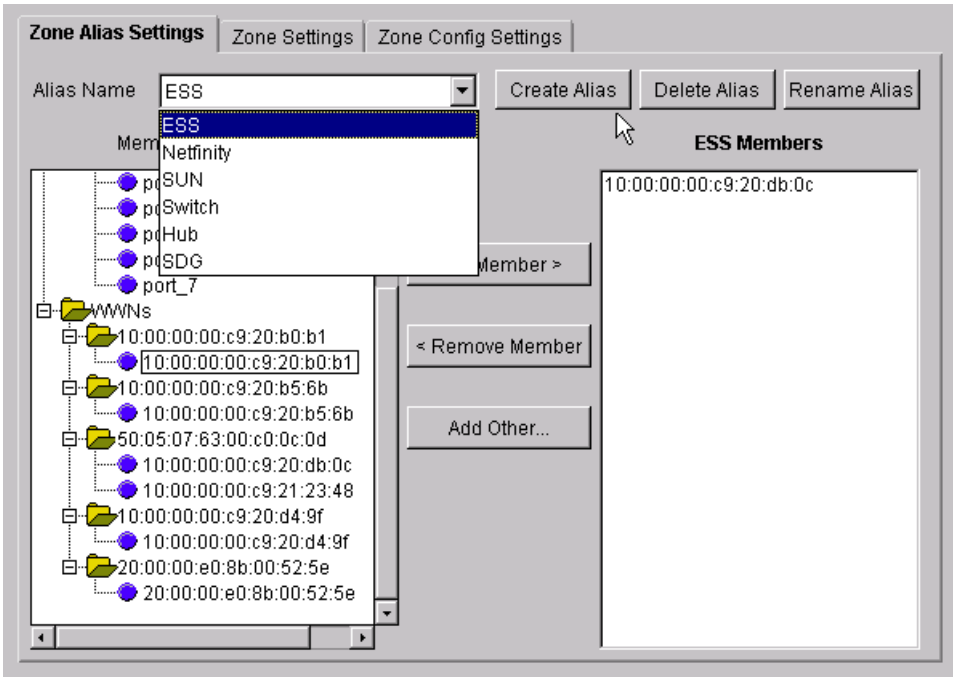

Figure 139. Zone Alias create example

6.4.2.4 Software zoning

As we have stated previously, one disadvantage of hardware zoning is that devices have to be connected to a specific port. If the device is connected to a different port which is outside of the zone, the whole zoning configuration will become unusable.

The members of software zoning can be defined as either a node WWN or a port WWN. They are unique.

With software zoning, your zoning configuration is not affected if the device is moved off a specific port. The WWNN or WWPN adds the consistency that port zoning does not offer. This is because the zone will follow the WWN. Software zoning also allows for the creation of symbolic names for the zone members.

An example of software zoning was illustrated when we defined the members of our NTzone (step 4 on page 165). In our example, 40:05:07:63:00:c0:0c:5d is a software zoning member.

Figure 140 shows this NTzone again.

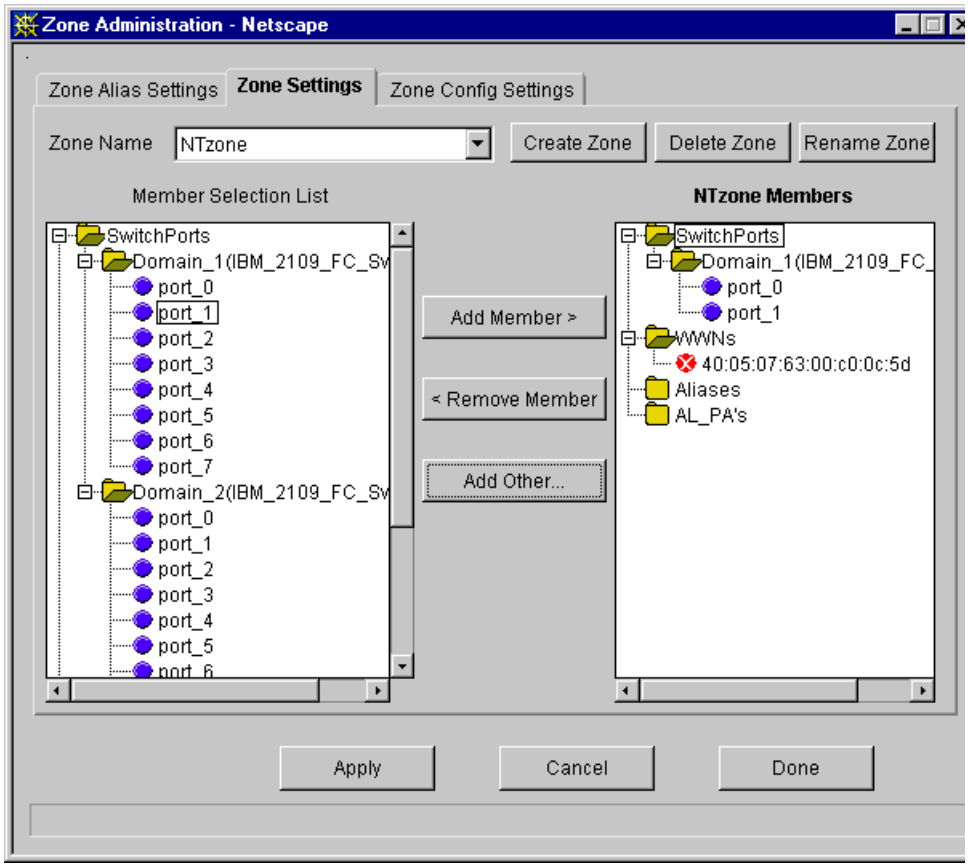

Figure 140. Software zoning representation

This process can be easily repeated for as many WWNN or WWPN that you require to be in a zone.

6.5 Cascading IBM 2109 switches

As your SAN begins to grow, you may want to add more switches. At some stage you may want to connect them together. This is known as cascading. In the following sections, we cover these topics:

- 6.5.1, “Cascading overview” on page 172
- 6.5.2, “Cascading examples” on page 172, showing cascaded fabric.
- 6.5.3, “Cascading using IBM switches” on page 174, showing 2109 capabilities.

- 6.5.4, “Implementing cascading with the IBM 2109-S08 switch” on page 175

6.5.1 Cascading overview

We have already stated that switches are among the core elements of a SAN fabric.

When you start your SAN business you might begin working with simple SAN fabric with perhaps only one switch and a few ports. But as you start implementing more and more host and storage devices in your SAN, you will reach limits where you have to add either more ports, or expand your fabric by adding more switches.

Interconnecting your first switch with newly added or other switches allows you to build much larger fabrics. Expanding the fabric in this manner is called *cascading*.

Cascading is a reliable, scalable, and cost effective way to build fabrics with hundreds or thousands of ports.

Cascading of switches provides the following benefits to a SAN environment:

- The fabric can be seamlessly extended. Additional switches can be added to the fabric, without powering down the existing fabric.
- You can easily increase the distance between various SAN participants.
- By adding more switches to the fabric, you increase connectivity by providing more available ports.
- Cascading provides high resilience in the fabric.
- With inter switch links (ISL) you can increase the bandwidth. The frames between the switches are delivered over all available data paths. So, the more ISLs you create, the faster the frame delivery will be. However careful consideration must be employed to ensure that a bottleneck is not introduced.
- When the fabric grows, the SNS is fully distributed across all the switches in fabric.
- With cascading you also provide greater fault tolerance within the fabric.

6.5.2 Cascading examples

A multiswitch fabric offers more flexibility to build a fault tolerant system.

The sophisticated path selection protocol allows multiple links to operate at the same time between any two switches in the fabric, so multiple redundant paths can be created. The extra links do not have to be reserved and maintained idle — all links carry traffic.

If a link goes down, the traffic that was carried over that link will be simply transferred to the other link(s). This transfer takes places automatically, with no human intervention, and in a very short time.

Even if a switch goes down, all other switches in the fabric and the end nodes connected to them are not affected.

Of course, if a switch goes down, all nodes connected to it will not be able to talk to each other, or to a node connected to another switch. However, a node running a mission critical application can have multiple Fibre Channel interfaces, each one connected to a different switch in the fabric, to overcome this problem.

This is not possible in a single switch fabric. In Figure 141, we show a six switch implementation of a fault tolerant fabric.

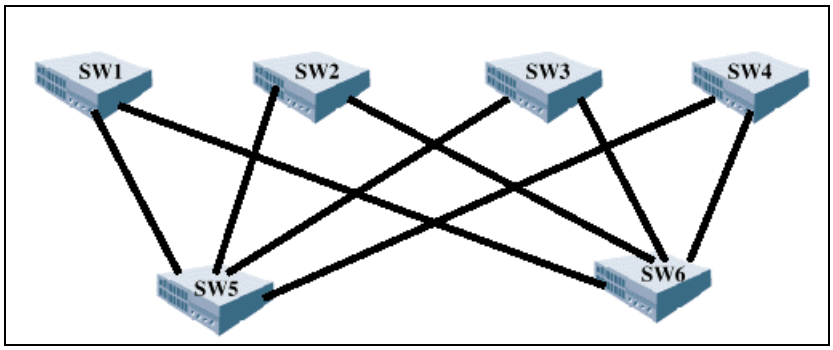

Figure 141. A fault tolerant fabric with six switches

This is an example of a completely fault tolerant fabric. If any link goes down, every switch is still able to communicate with all the other switches. If a switch goes down, all the remaining switches are able to communicate with each other.

Another example of a fault tolerant fabric with shorter distances between the switches is shown in Figure 142.

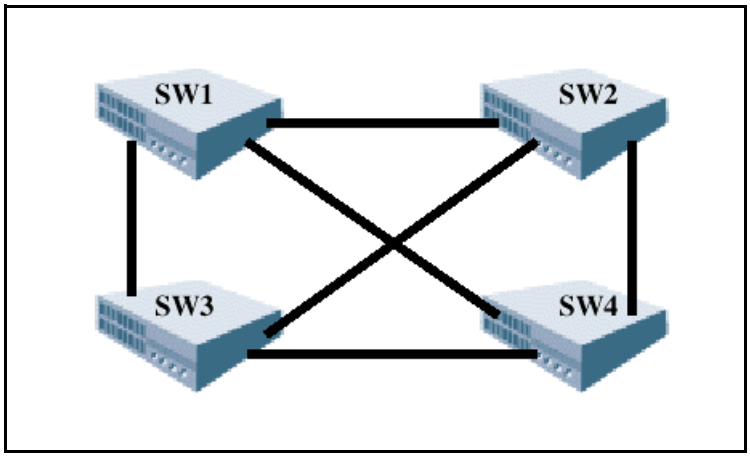

Figure 142. A fault tolerant fabric with four switches

Fabric fault tolerance, distance, or performance considerations determine the design of a multiswitch fabric; the number of involved devices decide the number of switches.

6.5.3 Cascading using IBM switches

All IBM 2109 type switches can be interconnected, and due to the switch software and embedded protocols, this is a non-disruptive topology change.

The 2109 ports are Fibre Channel G_Ports or FL_Ports. G_Ports can connect to either an external device (F_Port mode) or to another switch (E_Port mode). FL_Ports allow a Fibre Channel Arbitrated Loop to be connected to a (possibly cascaded) fabric, supporting communication between loop devices and fabric attached devices.

The IBM switch software automatically selects the port's mode of operation, so that no configuration is necessary.

This feature is extremely useful in a cascaded environment, because it allows a user to allocate more bandwidth between switches, if necessary, without having to plug in a new card, or even to change a configuration parameter. Even a port previously connected to an external device, such as a workstation, can be connected to another switch by simply unplugging the old cable and plugging in the new one, connected to another switch.

The 2109 software will detect the state transition and automatically configure the port to be an E_Port.

Theoretically, up to 239 switches are supported in an IBM SAN fabric. This allows thousands of fabric and Fibre Channel Arbitrated Loop connections. Practical implementations tend not to have taken advantage of this number as yet.

Considering the complexity and performance that cascading may introduce, the recommendation is to limit the number of cascaded switches to seven.

It should be noted that in a fabric environment some of the ports are used for inter switch links (ISL), that is E_Ports. This means that the maximum number of fabric attached device connections is actually less than the total number of ports.

The number of ports used as ISLs depends on the fabric configuration. This is dictated by the amount of bandwidth required between switches, and by the number of desired redundant links, to maintain connectivity in the case of link or switch failures.

6.5.4 Implementing cascading with the IBM 2109-S08 switch

We will use our basic scenario with one host, a 2109_S08 switch and an ESS storage device and add a new switch to the configuration.

6.5.4.1 Adding a new switch

In 6.2.1, “Installing the 2109-S08 switch” on page 143, we have described the process to successfully implement a new switch. A new switch is a switch that has not been previously connected to a fabric with zoning configured, and which has had no zone configuration data entered into it. To return a switch to new switch state, you must use the `cfgClear` command on the switch that you are adding, before connecting it to the fabric. The `cfgClear` command must be used with caution as it removes all zone information from the fabric. If you are unsure of the consequence of this, do not issue the command.

We show an alternative method to add a previously configured switch in 6.5.4.2, “Adding a switch that has been previously configured” on page 176.

To implement the new switch, make sure you follow these steps:

1. Bring in the new switch.
2. Establish one or more paths to the existing switch by connecting the equivalent ports using fiber-optic cables.
3. Check the switch parameter settings.

You can find a list of all the possible switch parameters in:

- *IBM 2109 Model S08 Installation and Service Guide*, SC26-7350
- *IBM 2109 Model S08 User's Guide*, SC26-7349

To change any or all of these settings, you must Telnet to the 2109 Switch and enter the `configure` command.

4. Power on the new switch.
5. The IBM 2109 switch software detects the fabric expansion and automatically configures the related ports as E-Ports.
6. As an administrator you can see both switches in the fabric.
7. The port LEDs on the switch are solid green.
8. The extended fabric is working.

6.5.4.2 Adding a switch that has been previously configured

The new switch process works fine under the assumption that all parameters relating to the switches in the extended fabric are set properly and match each other.

However, this will not be the case if the switch to be added has been used before and has specific parameter settings.

As a result, the interconnecting port LED will show a slow, flashing green LED on your Switch View panel, like that shown in Figure 143.

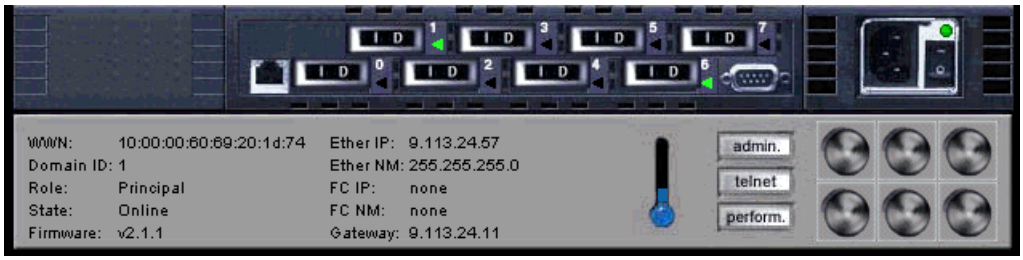

Figure 143. Switch view panel

The slow LED means the ports are online, but are segmented and cannot establish a fabric. The error message says:

```
Please check for loopback cable or incompatible switch
```

From the administrators Fabric View panel, you will see only one switch and not the complete fabric as depicted in Figure 144.

.

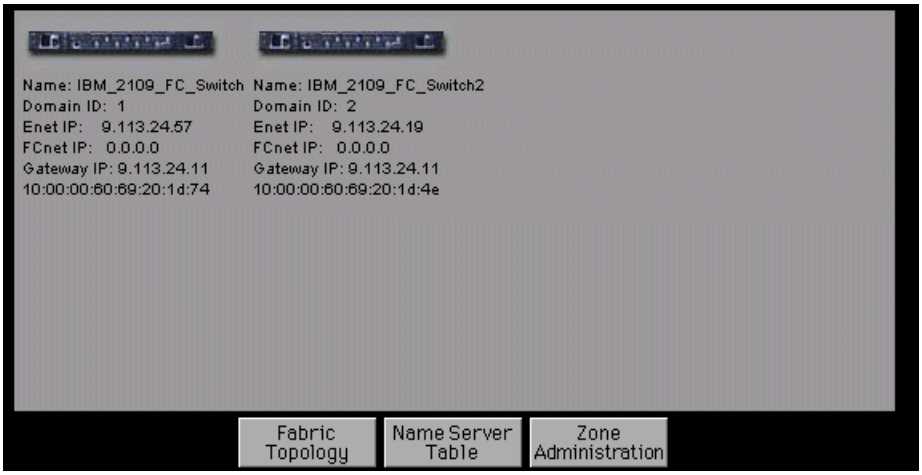

Figure 144. Fabric view panel

To make the switch known in the fabric you must take the following actions:

1. Telnet to the switch that you are adding, for example `telnet 9.113.24.57` and press Enter.
2. Login, enter your userid and password, issue the command `switchdisable`. and press Enter. This is necessary to modify switch parameters.
3. Issue the `configdefault` command and press Enter to set the switch's configuration parameters back to the factory default settings. Any zone information that was in this switch will be deleted.
4. The `configdefault` command is navigated by entering a series of collapsible top-level menus. However, there is another alternative that you can try before issuing the `configdefault` command. If you know the configuration parameters that relate to your environment, you can enter the `configure` command. Then:
   - Each menu divides up the various switch configuration parameters into logical groupings. These are fabric parameter, virtual channel parameters, arbitrated loop parameters, and system service parameters. Enter the parameters that reflect your environment.
5. Issue the `switchenable` command.
6. Issue the `reboot` command.

   After >switch:admin> **configDefault** the switch will answer with: `committing configuration...done`.

7. Make sure the two switches are working in different domains. This is mandatory.

The complete `configdefault` command sequence is shown in Figure 145.

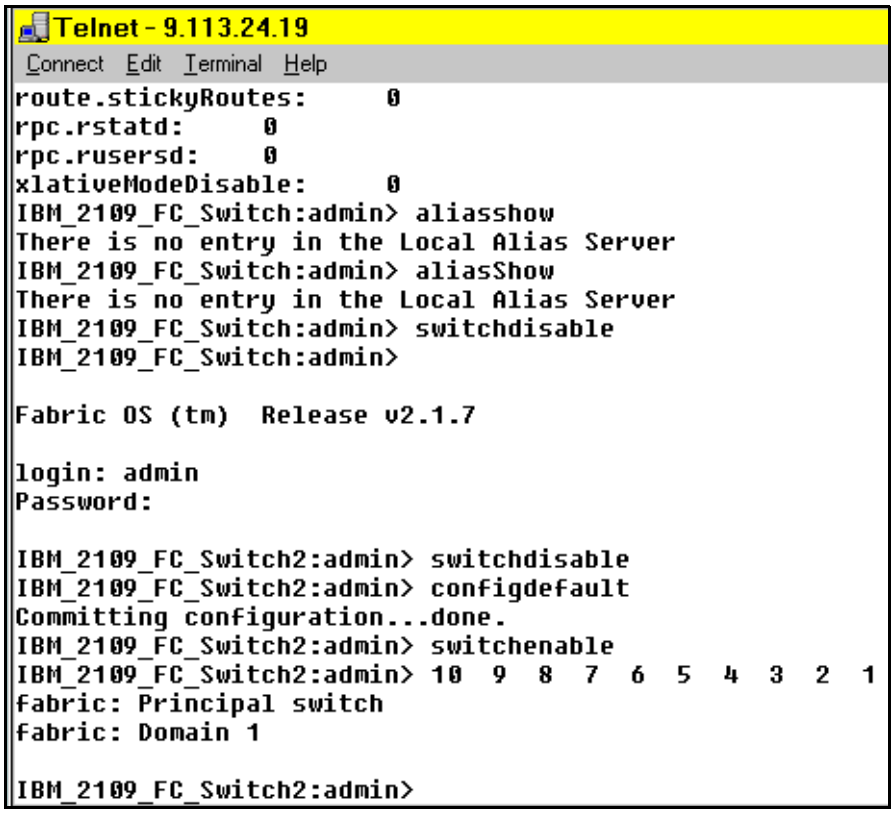

Figure 145. Setting default switch configuration

Should you choose to enter your own configuration parameters, be aware that these can significantly affect switch performance. The default factory settings should not be changed unless you are an experienced SAN administrator, and you have recorded the original settings before you make the changes.

The `configDefault` command is used to reset some of the switch configuration values to their factory default values.

In addition, this command configures the switch to boot from its internal firmware if it has been previously configured to boot from the network. This command may not be executed on an enabled switch; you must first disable the switch using the `switchDisable` command.

Because some configuration parameters are cached by the switch, it is recommended that the switch is rebooted immediately following the execution of the `configDefault`; otherwise, unexpected behavior may result.

With the exception of the following, all configuration parameters are reset to their default values:

- World Wide Name
- Ethernet MAC address
- Ethernet IP address and subnetmask
- IP Gateway address
- OEM customization
- SNMP configuration
- Zoning configuration
- System name

Now, cascading in our target SAN Fabric scenario will work. On the Fabric View panel, you will see both switches, as shown in Figure 146, and all switch port LEDs are showing a steady, solid green.

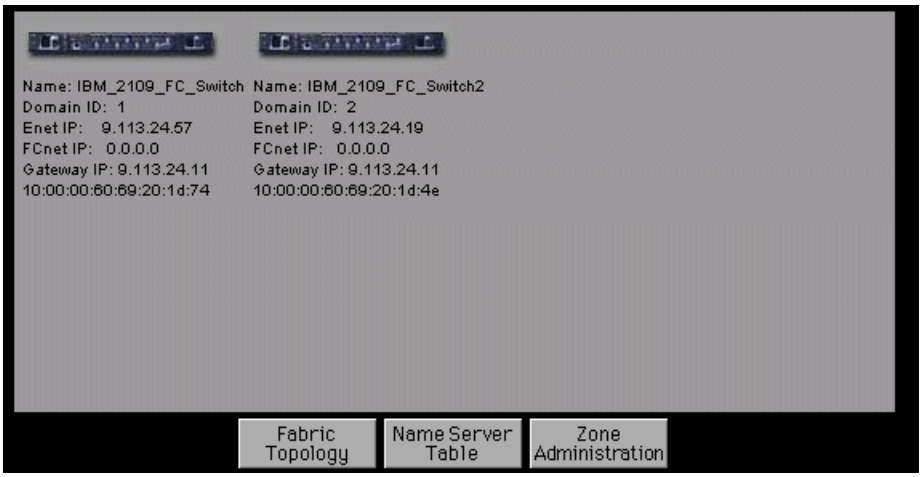

Figure 146. Cascaded fabric view

We have now successfully implemented a cascaded fabric.

6.6 Related information

IBM 2109 S08 Switch Service Guide, SC26-7350

- This guide introduces the product and lists the features you can order. It also provides steps for the installation, configuration and servicing of the IBM 2109 S08 Switch.

IBM 2109 S16 Switch User's Guide, SC26-7351

- This guide introduces the IBM 2109 S16 switch and its features. It also provides information on using the IBM StorWatch SAN Fibre Channel Switch Specialist, setting up zoning, and methods for managing the IBM 2109 S16 switch remotely.

IBM 2109 S16 Switch Service Guide, SC26-7352

- This guide introduces the product and lists the features you can order. It also provides steps for the installation, configuration, and servicing of the IBM 2109 S16 switch.

Translated External Devices/Safety Information, SA26-7003

- This book provides translations of the danger and caution notices used in IBM 2109 Switch publications.

Electrical Safety for IBM Customer Engineers, S229-8124

To get specific details on models and firmware that the switch supports, refer to this Web site:

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

Chapter 7. Implementing the McDATA ED-5000

Since the early 1990's, McDATA has provided IBM with fiber switches for the S/390 world, the ESCON directors. They have built up experience in creating the highly available and reliable equipment necessary to meet the demands of S/390 environment. Any SAN component needs to be as reliable and as available as the large systems to which they are attached.

With the McDATA Enterprise Director ED-5000 (ED-5000) and its Fibre Channel interface, McDATA has moved this proven technology towards the open market. With open systems closing up on S/390 systems, choosing the McDATA switch allows for enterprise class availability and performance that is usually found in mission critical data centers.

In this chapter, we show how to install and manage a McDATA ED-5000. This device is ordered through IBM as the McDATA Enterprise Fibre Channel Director, 2032-001.

The topics covered in this chapter are:

7.1 Our ITSO environment and the installation steps

We configured the ESS as described in Chapter 4, "Configuring the ESS with native Fibre Channel" on page 63, to provide access through the McDATA ED-5000 to the host. Our configuration is shown in Figure 147. This setup will be the starting point for our McDATA SAN.

In the following sections, we cover these topics:

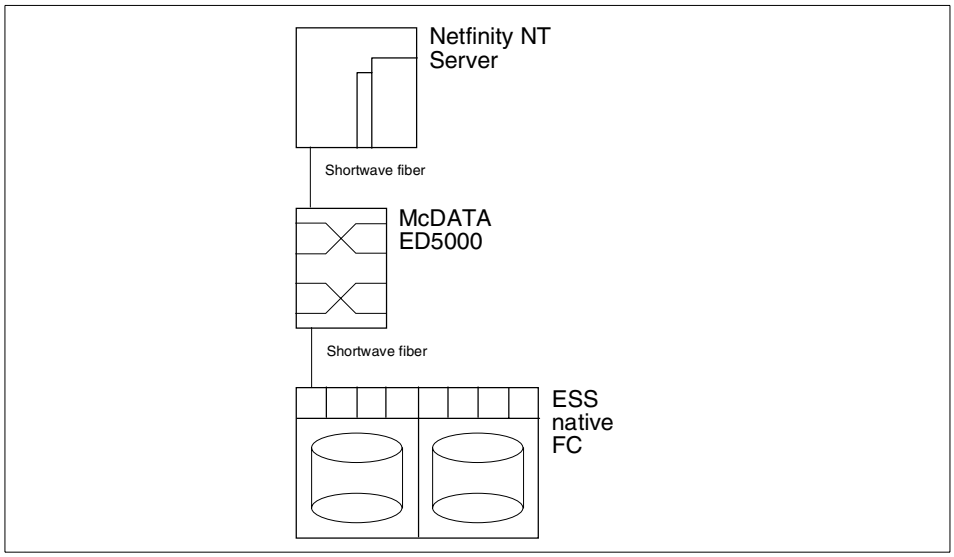

Figure 147. Simple McDATA switch setup

7.1.1 ED-5000 overview

At the heart of our McDATA SAN is the ED-5000. The ED-5000 is a Fibre Channel switch with a maximum of 32 Generic Ports (G_Port). These ports, when connected to another ED-5000, act as Expansion Ports (E_Ports) to setup inter switch links (ISL). Connected to hosts or storage devices, they act as fabric ports (F_Port). The ports are on cards with four ports each. Most parts of the switch are hot pluggable field replacable units (FRU). The design is geared towards high reliability and availability. The switch supports features like internal failover to redundant FRUs, and fault detection. The maintenance firmware loads and upgrades can be done without downtime and a call-home feature is also provided. The management and monitoring software does not run directly on the switch. It is provided through an Ethernet connected PC which is called the Enterprise Fabric Connectivity (EFC) server. This PC is shipped as part of the ED-5000 package.

For in-depth information on the McDATA, also see the McDATA ED-5000 manuals provided with the switch:

- *ED-5000 Enterprise Fibre Channel Director Planning Manual*
- *Enterprise Fabric Connectivity Manager User Manual*
- *ED-5000 Enterprise Fibre Channel Director User Manual*
- *ED-5000 Enterprise Fibre Channel Director Installation Manual*
- *ED-5000 Enterprise Fibre Channel Director Service Manual*

- *ED-5000 Product Manager User Reference Guide*
- *EFC and Fabric Manager User Reference Guide*

7.1.2 Hardware overview

The hardware that we will use in this SAN is:

- 2 McDATA Enterprise Director ED 5000
  - 32 ports
  - 50 micron shortwave cable
  - 9 micron longwave cable
- 1 Netfinity NT Server 5500
  - 1 FC HBA Emulex 7000
- 1 Netfinity NT Server 5500
  - 1 FC HBA QLogic 2200
- 1 ESS 2105-F20 with Feature Code 3022 (native FC adapter)
  - 10 FC adapter, 10 ports
  - 1 SCSI adapter, 2 ports
- 3 RS/6000 AIX servers
  - 1 FC HBA Emulex 7000

7.1.3 Software overview

The software that we use in this SAN is:

- Netfinity: Microsoft Windows NT 4.0 Service Pack 5
- RS/6000: AIX
- Emulex 7000
- QLogic 2200
- ESS: Microcode (1.1)

7.1.4 ED-5000 installation steps

The installation of the switch and the EFC server is straight forward. The ED-5000 ships with all necessary components. The McDATA SAN was implemented using the following steps; we have included them for reference. These steps are also provided with the *ED-5000 Enterprise Fibre Channel Director Installation Manual,* which should be referred to if more detailed information is required.

In 7.2, “Setting up an environment for using and managing a McDATA SAN” on page 187, we show how to manage this environment once it is installed. If you have an environment with the McDATA ED-5000 already installed, you may skip this section.

We do not carry out every task, for example, setting up for SNMP; nor do we describe every single step in its entirety. However, we show how we built and managed our McDATA fabric in a task-oriented way.

1. Verify installation requirements.

   We know what equipment we want to use and how we want to configure it. In our case, we want to first build a single switch fabric with one ESS and one Windows NT server. Later on, we want to extend the setup with more hosts, for example, RS/6000. Once this is established, we build zones and add another ED-5000 to build a cascaded fabric. The hosts are already attached in our environment and so is the ESS. What we need to ensure is that everything relating to the physical and logical installation of the ED-5000 is complete. For instance, we need to make sure we have access to power supplies, and access to the local network. What we need to make sure, before attempting any installation, is that the necessary resources, whatever they may be, are available to support us or can be scheduled to become available.

   In Appendix A of the *ED-5000 Enterprise Fibre Channel Director Planning Manual*', there is a checklist which is useful for keeping track of the project and its progress.

2. Determine local area network addresses.

   Note: If you plan to connect both the EFC Server and the switch on a dedicated LAN, you do not have to change the network information for the first switch. Of course, if you later add more switches, you will need to change the network information for the new devices.

   This is referred to in 7.2.1, “Setting up the network environment” on page 187.

3. Install and configure the Ethernet hub.

   We will not be installing the optional Ethernet hub. This is described in the *ED-5000 Enterprise Fibre Channel Director Installation Manual*, should you want to implement this.

4. Unpack and inspect the ED-5000.

   Note: Verify that your ED-5000 is not damaged and that no items are missing. For a complete list of items that are shipped, refer to the *ED-5000 Enterprise Fibre Channel Director Installation Manual*'.

5. Position the ED-5000 and route power cables.

   This consists of placing the ED-5000 in the planned area with enough ventilation and then connecting the power and the network. Note: You may also have to adjust the leveling pads.

6. Switch power on to the ED-5000 and verify operation.

   After the hardware installation, we power on the ED-5000. After power on, the ED-5000 runs a couple of Power-On Self-Tests (POST). At the successful completion of the tests, the power LED is illuminated and the operator panel shows that the ED-5000 is online.

7. Set IP address, Subnetmask, and Gateway address.

   To configure the ED-5000 from within the EFC Manager, the switch itself needs to be configured for the network. For operations like this, we will use the operator panel at the front of the switch. We can display and change the status of the switch and the current network information there. Since the switch is initially set up to work on the dedicated Ethernet segment, we need to update the network information that was determined in step 2. Refer to 7.2.1, “Setting up the network environment” on page 187 for more information.

8. Install the ED-5000 on the LAN.

   We configured the switch as a member of the LAN and now we are going to connect it. We use the supplied Ethernet cable and connect it to a Ethernet hub. Then we route it to the cable access hole inside of the ED-5000. Going through the cable restraint mechanism, we connect the cable to the left Ethernet plug. We do the same with the second Ethernet cable on the backup card to the left.

9. Unpack and inspect the EFC server.

   Note: As with the ED-5000 itself, verify that the EFC Server is not damaged and that all parts are shipped. For a complete list of items to be shipped refer to the *ED-5000 Enterprise Fibre Channel Director Installation Manual*.

10. Install the EFC server on the LAN.

    This consists in plugging together the PC and the peripherals shipped to set up the EFC server, for example, connecting the monitor, the keyboard, and the mouse to the PC. We connect the inboard Ethernet interface to our LAN. Note: Don’t forget to connect the power supply.

11. Configure EFC server network information.

    Note: Use the information determined in step 2 to install the EFC server in the LAN. To log on to Windows NT, use the user id `'Administrator'` and the password `'password'`, both of which are case sensitive.

12. Set date and time on the EFC server.

    Note: As the audit and event logs of the EFC Manager are time stamped with the date and time of the EFC Server, make sure that the time and date on the EFC Server are correct.

13. Configure Windows NT users on the EFC server.

    Note: If you plan to allow more users to access the EFC Server, then we recommend that you create Windows NT users on the EFC server and change the Administrator password.

14. Configure the call-home feature.

    Note: The EFC Server has a built in modem which you can configure to report ED-5000 problems to remote service support facility.

15. Assign EFC Manager user names and passwords.

    The EFC Manager has the ability to manage different users with different rights. Refer to 7.2.3, “Defining users on the EFC Manager” on page 192 for more information.

16. Identify the ED-5000 to the EFC Manager.

    The ED-5000 has to be made known to the EFC Manager if it is to be managed from within the EFC Manager. This is described in 7.3, “Managing the ED-5000 fabric” on page 203.

17. Record information for EFC Server restore.

    Note: You will need to record some information in case you have to restore the EFC manager software.

18. Verify EFC Server communication with the ED-5000.

    The operational state of the devices and the communication is indicated with different colors and symbols in the EFC Manager.

19. Set date and time on the ED-5000.

    The logs of the Product Manager are stamped with the time and date of the ED-5000. Note: To be able to use the logs efficiently, set the time using the Product Manager.

20. Test remote notification.

    Note: If you want to be sure that remote notification works as expected, test it using the Maintenance icon in the Product View of the EFC Manager.

21. Configure the EFC Manager and ED-5000 Product Manager.

    For example, this consists of configuring the identification, ports, and the operating parameters.

22. Back up the configuration.

    The EFC Server ships with an internal Zip drive. It is used to mirror the EFC Manager data every time it has changed.

23. Cable the ED-5000 device ports.

    Last, but not least, we connect the Fibre Channel cables to the Fibre Channel cards of the ED-5000.

24. Configuring a multiswitch fabric.

    This consists of connecting two switches together using a fiber-optic connection known as an inter switch link. Refer to 7.5, “Cascading with McDATA - building a multiswitch fabric” on page 237.

7.2 Setting up an environment for using and managing a McDATA SAN

Now that we have installed the ED-5000 in a base configuration, in the following sections, we cover these topics:

- 7.2.1, “Setting up the network environment” on page 187
- 7.2.2, “Logging on to the EFC server and the EFC Manager” on page 190
- 7.2.3, “Defining users on the EFC Manager” on page 192
- 7.2.4, “Installing the EFC Manager on remote workstations” on page 196

7.2.1 Setting up the network environment

First, we need to think about how we want to integrate the McDATA in our network environment. Do we want the EFC server accessible through our corporate intranet or do we only want to access it locally? Second, do we want to establish a dedicated network segment to connect the switch with the EFC server, or do we want to use the existing LAN for this?

7.2.1.1 Our network setup

We want to be able to reach the EFC server through the network, which is also reachable from remote using IP, so we need to install the EFC server in the LAN.

Also, we chose to install the McDATA ED-5000 in our laboratory network as well. This configuration is shown in Figure 148.

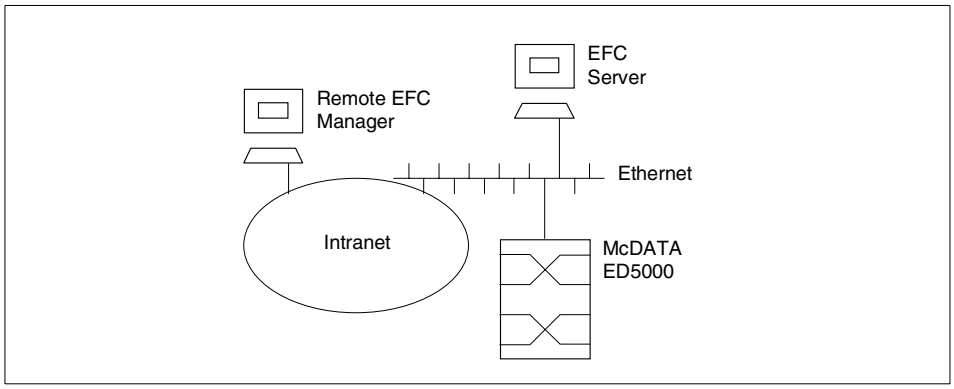

Figure 148. Our ED-5000 setup

To set up the switch and the EFC server, we used the following network information:

For the EFC server:

- IP address: 9.113.25.62

  The unique address of the EFC Server in the corporate intranet. We use it with a Web browser to install the EFC manager on a remote workstation and to access the EFC Server from the remote workstation.

- Subnetmask: 255.255.254.0

  Our intranet is divided into several subnets. The Subnetmask is used to distinguish between the net part and the host part of the IP address. The host needs this information to know how to treat the addresses it uses.

- Domain Name System (DNS) host name: EFC-Server-1

  Numbers, for example, IP addresses, are sometimes difficult for people to remember. Therefore, a service called the Domain Name System is used to assign host names to IP addresses, and it makes the destination transparent and reachable with an easy to remember host name.

- DNS domain name: sanjose.ibm.com

  The name says that the DNS is using a hierarchy of domains for addressing. The host name itself is not unique; to be able to reach the hosts given, we need a fully qualified host name which consists of the host name and the domain it is in.

- Gateway address: 9.113.42.250

  This is the default router which has to be known to the IP stack to route to a destination outside the current LAN using IP.

For information about how to set up the network on the Windows NT PC, refer to the *ED-5000 Enterprise Fibre Channel Director Installation Manual*.

The director needs to have following information too, to be accessible from the management software.

For ED-5000-1, this is:

- IP address: 9.113.25.63
- Subnet mask: 255.255.254.0
- Gateway address: 9.113.42.250

To set the network information on the switch, we use the operator panel. We can scroll though the information displayed on the panel with the Advance button. To change specific information, we display it using this button. As we want to change it, we press the Clear button to get a cursor in the entry field. With the Detail button, we increment the value of the field we have activated. Pressing the Clear button again gets us to the next field to change. If all entries are correct, we use the Entry button to save the changes. This has to be done with the IP address, the Subnetmask, and the Gateway address.

7.2.1.2 Suggested network setup

We suggest, for security and availability reasons, using the optional Ethernet hub to establish a dedicated Ethernet segment to connect the EFC server and the switch, which should look similar to that shown in Figure 149.

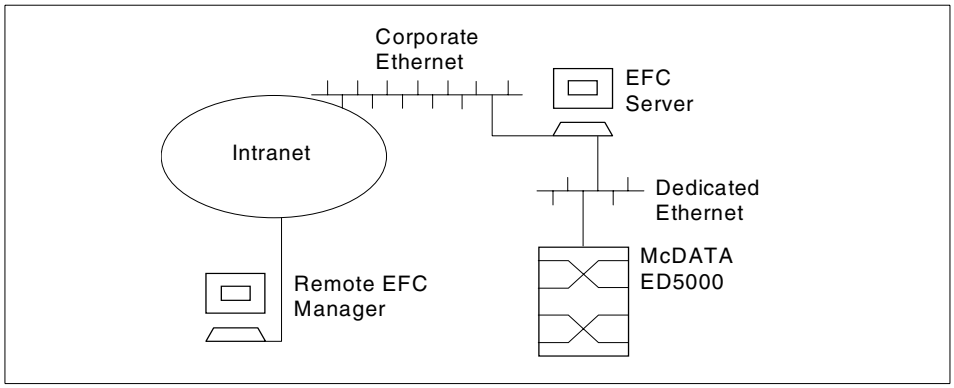

Figure 149. Suggested ED-5000 setup

If you plan to use the dedicated Ethernet between the EFC server and the ED-5000, then you only need to set up the EFC server for IP.

7.2.2 Logging on to the EFC server and the EFC Manager

We logon to Windows NT with the user ID, `Administrator`, and the password, `password`, both of which are case sensitive. After logging on to Windows NT, we get the login window for the Enterprise Fabric Connectivity Manager (EFC Manager), which is installed and operational on the system.

The EFC Manager is a centralized tool for the McDATA network. It is a Java based GUI that shows a graphical view of all managed devices and functions as an entry point for all common management and monitoring tasks.

For detailed information about the EFC Manager and how to use it, refer to the *Enterprise Fabric Connectivity Manager User Manual*.

To start the EFC Manager, we log on with the same user and password we used when logging on to Windows NT, `Administrator` and `password`. We are locally working on the EFC server and, therefore, we specify in the EFC server entry field `localhost`, which is shown in Figure 150.

Figure 150. Logging in to the EFC Manager on the EFC Server

After logging on to the EFC Manager, it opens with the Product View, shown in Figure 151.

On the left, there is a panel that allows us to perform various configuration and monitoring tasks. The main window is empty, because there are no devices configured on the EFC Manager. The same applies to the Fabric View of the EFC Manager. We can switch to it by using the View button on the panel, as shown in Figure 151.

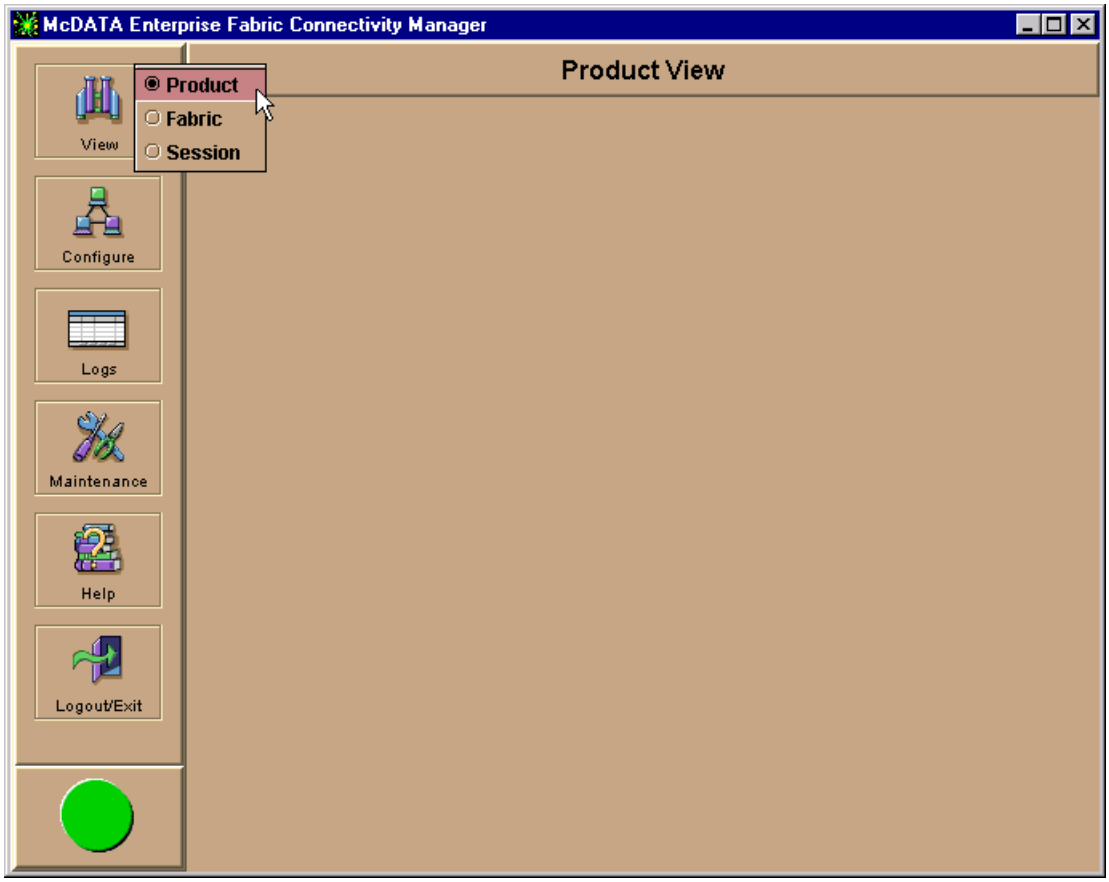

Figure 151. EFC Manager, product view with no devices configured

7.2.3 Defining users on the EFC Manager

We define users on the EFC Manager, because we do not want the Administrator user ID to be used remotely, so we create a new user and use that for remote access.

We can define up to 16 users for the EFC Manager, but only a maximum of four can log on concurrently. This includes the user of the EFC Manager running locally on the EFC server.

From the Product View, we go to **Configure -> Users** on the button panel, on the left side of the window as shown in Figure 152.

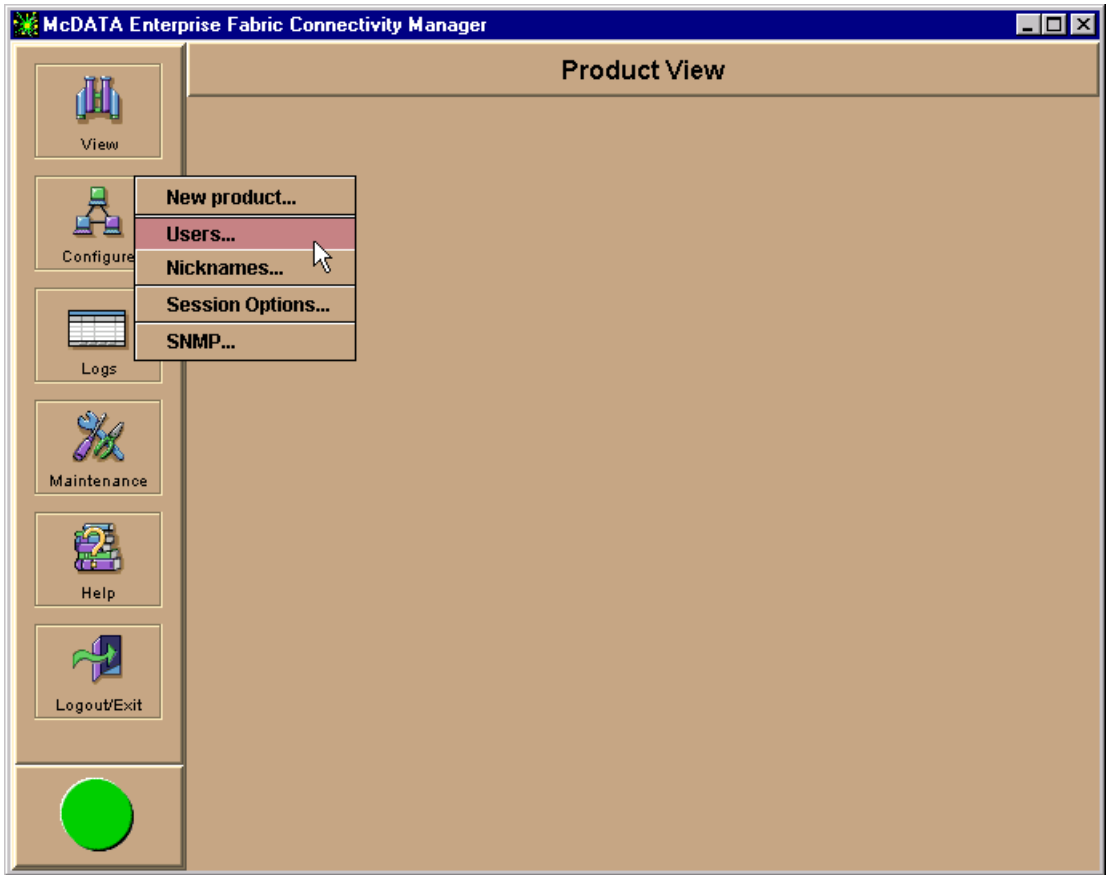

Figure 152. Configuring users

We are presented with a list of the defined users and the options to add users, modify existing users, view the rights of a user, and delete users. We will add another user by pressing the New button and then specifying the name, password, and description of the new user. Also, this window is used to specify the rights that the new user should have. This is shown in Figure 153.

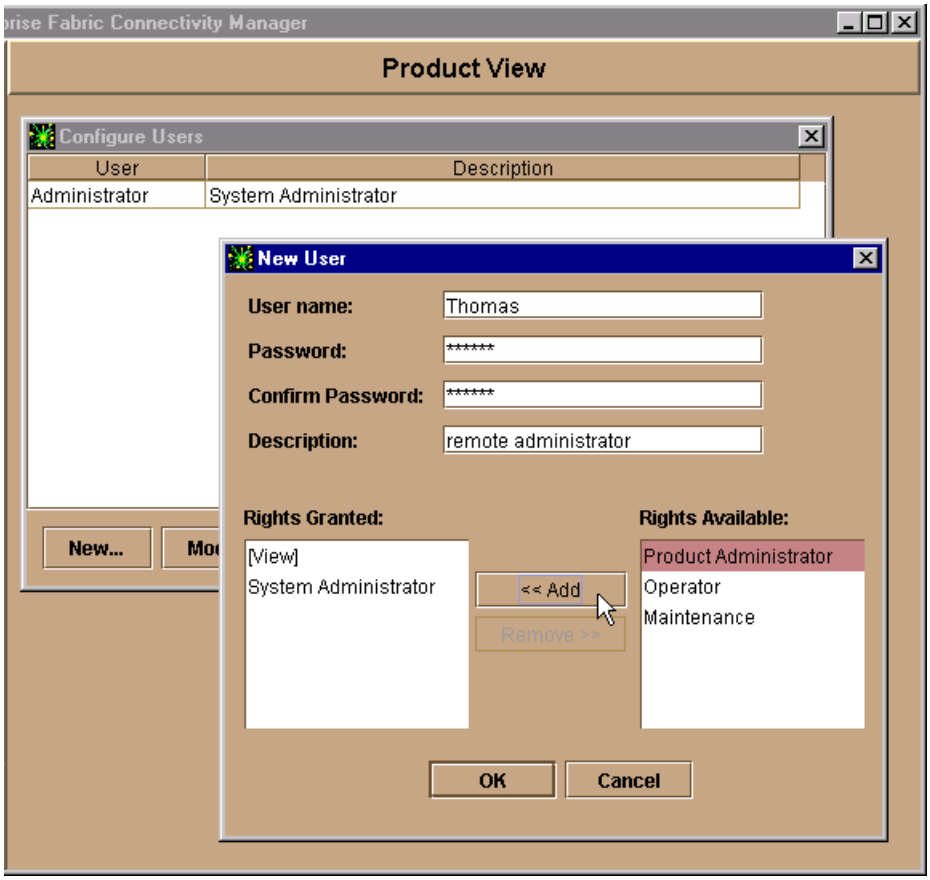

Figure 153. Defining new user

To assign rights to the user, we click on one of the Rights Available and the Add button. The rights are:

- System Administrator
- Product Administrator
- Operator
- Maintenance
- View

The *System Administrator* right grants access to every control and configuration task that needs to be performed and can be viewed as the highest level of authority. All new users initially have view rights and this cannot be removed. For a table of user rights of product manager functions, refer to the *Enterprise Fabric Connectivity Manager User Manual*.

To change the settings for a user, for instance, to change the password, we go to **Configure -> Users**. With the Modify button we get a window, similar to the New window, where we can change our password and the user rights. This is shown in Figure 154.

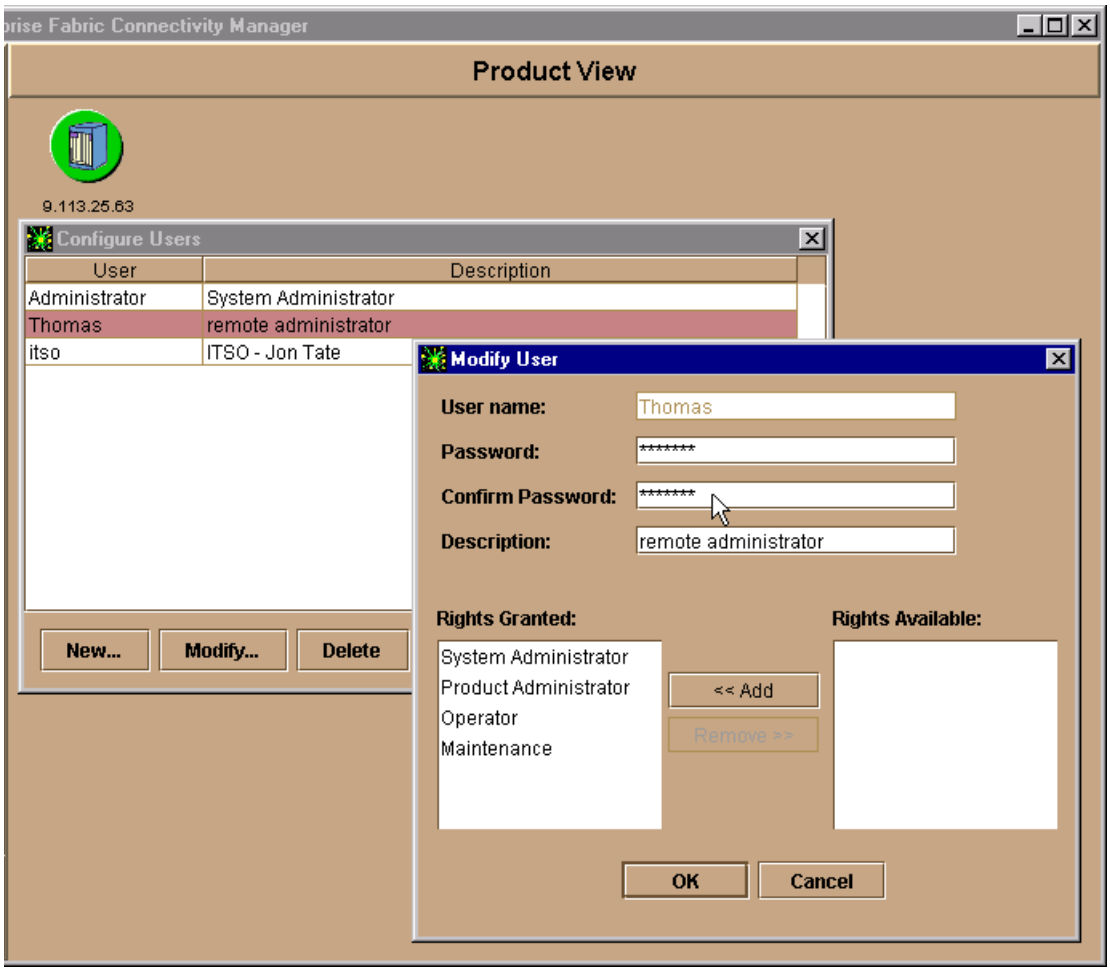

Figure 154. Modify users

Once the new user has been defined, we can login to the EFC server with the newly created user ID and password. To be able to connect from a remote workstation, the EFC Manager has to be installed on this workstation. This is described in 7.2.4, “Installing the EFC Manager on remote workstations” on page 196.

7.2.4 Installing the EFC Manager on remote workstations

So, we know that the EFC manager is running properly and we have defined users that can access it remotely. Now, we have to install the EFC Manager on a remote workstation. This can be a PC running Windows NT, a SUN or a HP workstation. In our case, this will be a PC running Windows NT.

To use the EFC Manager from a remote workstation, we need to download and install the code on our workstation. The code is downloaded from within the EFC server. In our case, it is an already installed EFC server of another ED-5000. The download and installation of the EFC Manager is done using a Web and Java based installation procedure. All we need is a Web browser, and we will use Netscape Navigator. In the Uniform Resource Locator field (URL) of the Navigator, we point to the address of the EFC server, 9.113.24.119, to access the initial page on the EFC server.

7.2.4.1 Changing the Netscape Navigator preferences

If you can access the page, but only the background is visible, it is likely that you may have to change the preferences of the Netscape Navigator. To accomplish this, go to the Preferences by selecting **Edit -> Preferences**. On the Advanced tab, uncheck the check box **Enable Style Sheets** as shown in Figure 155.

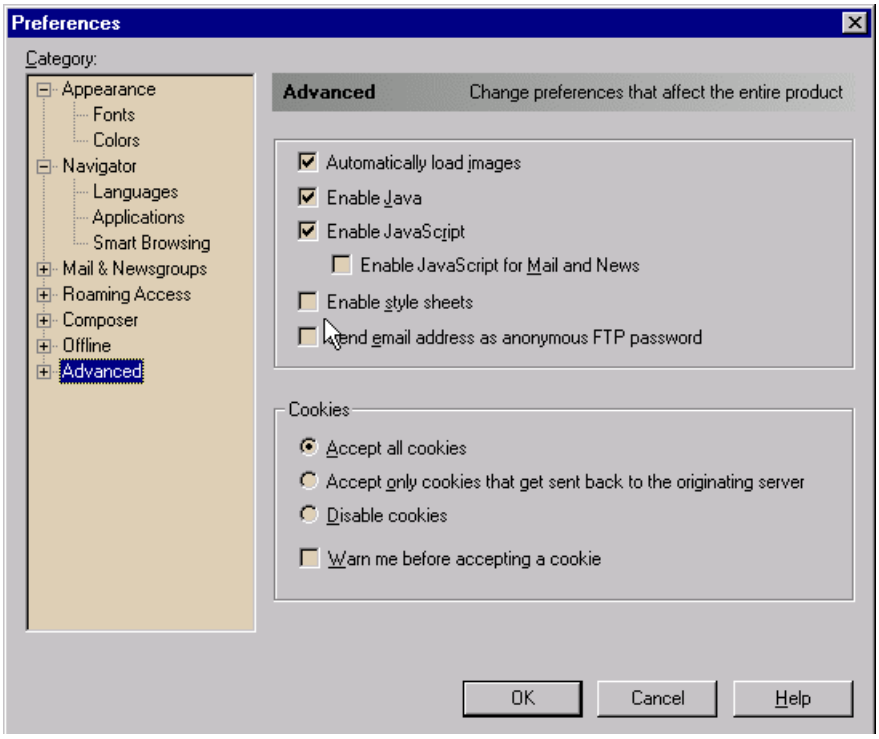

Figure 155. Netscape preferences, disabling style sheets

7.2.4.2 Downloading and installing the EFC Manager

This takes us to the start page for the remote EFC Manager installation, as shown in Figure 156, where we can choose the operating system of our remote workstation.

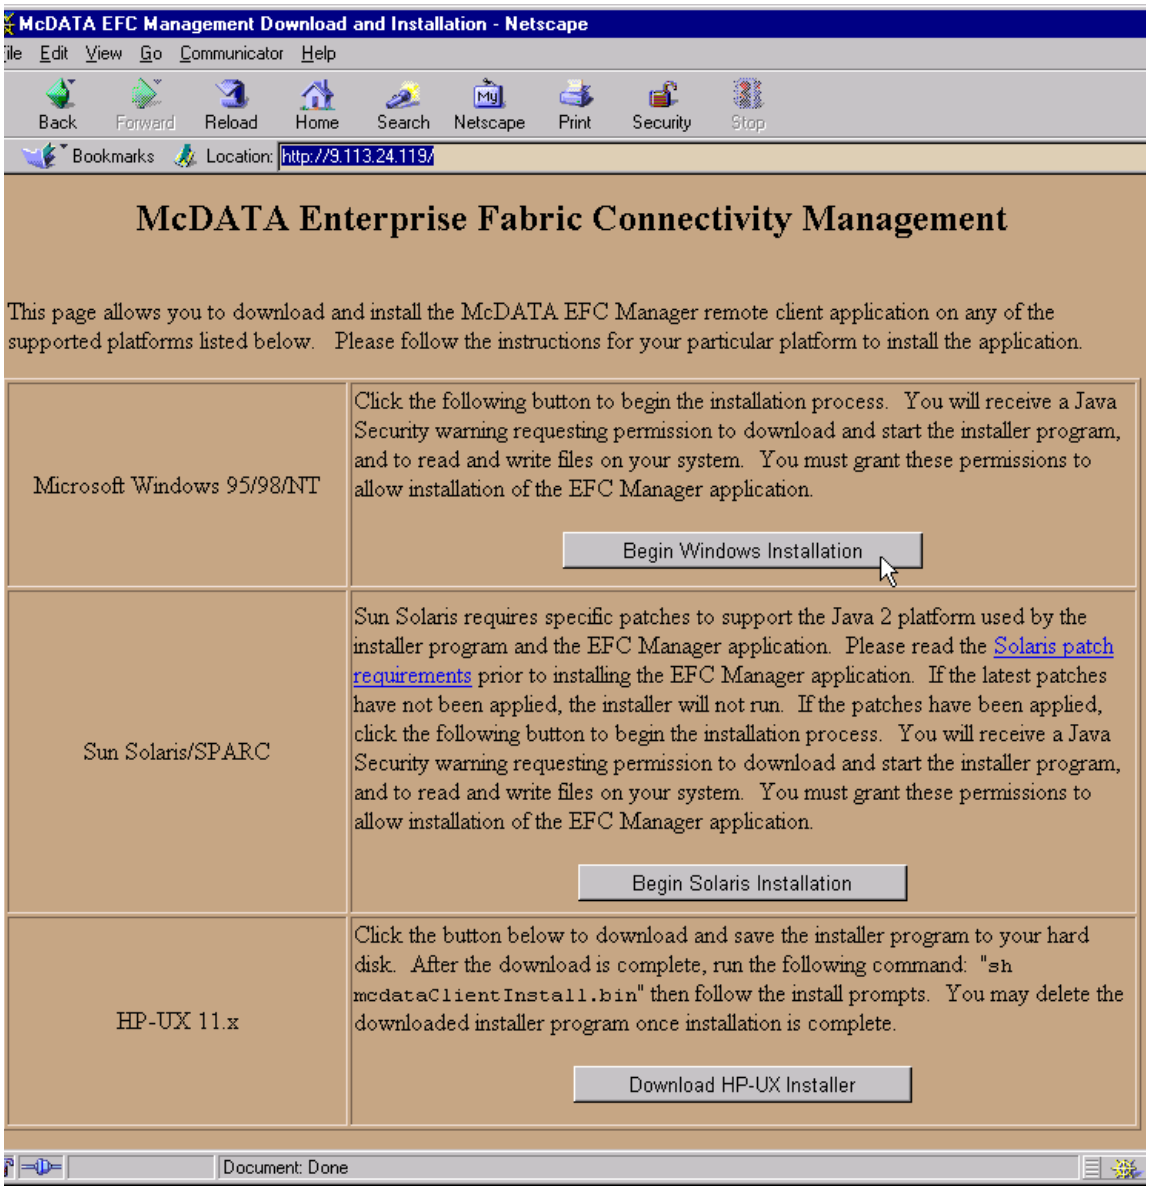

Figure 156. Start page for remote EFC Manager installation

We will install the EFC Manager software on a Microsoft Windows 95/98/NT system, so we select that option. After doing so, we are taken to the next Web page to start the download and installation procedure. Immediately, we are prompted to grant additional privileges to the Java based installation

software. First, we have to grant the right to start programs, and then grant the right to read, modify, or delete files. If we do not grant the additional privileges, we will not be able to perform the installation. The button to start the download is covered until we grant the rights, as shown in Figure 157.

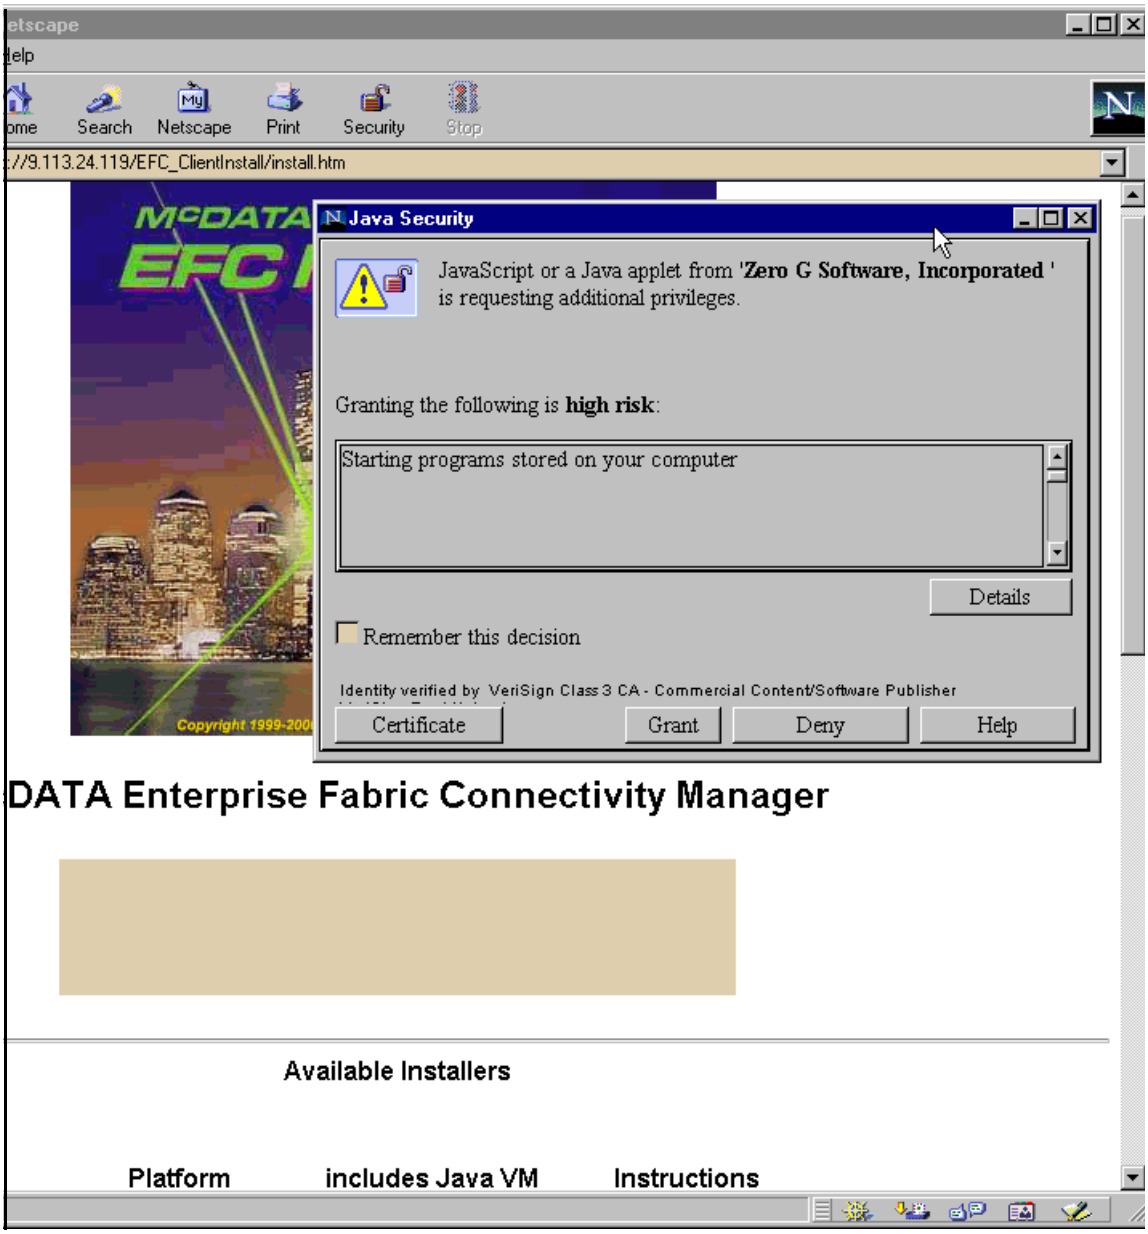

Figure 157. Granting additional rights to the EFC Manager installation software

Due to the fact that, for security reasons, a Java applet is not allowed to perform the tasks mentioned above, this warning message appears. Only allow Java programs to perform like this from trusted sources.

After this, the button to start the InstallAnywhere Web Installer appears, and we are able to start the installation, as shown in Figure 158.

Figure 158. Starting the installation of the EFC Manager

After pressing the button to start the installation process, the software starts downloading to our local machine and begins the installation. We now follow the instructions to install the EFC Manager. After confirming the licence agreement, we get information about which version we are going to install, as shown in Figure 159. The final step is to tell the installation program where to put the program files.

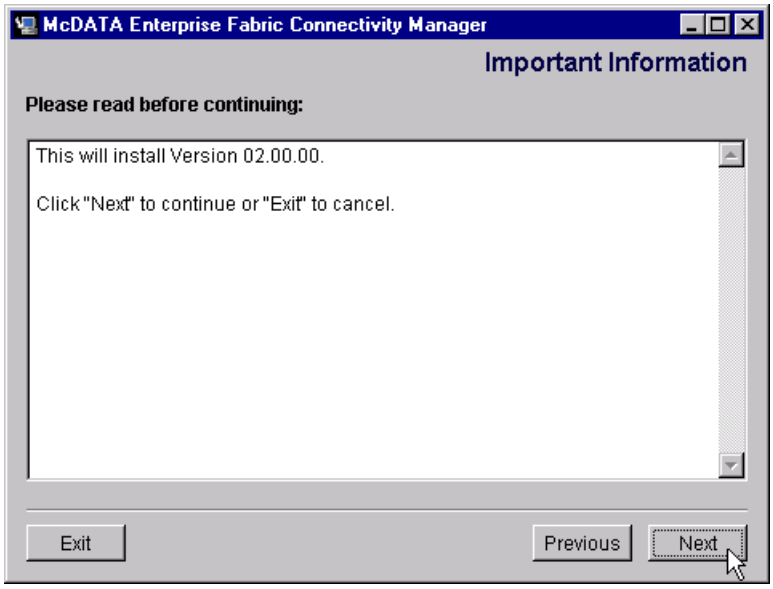

Figure 159. EFC Manager version

After finishing the installation, there will be a shortcut to the EFC Manager on the desktop, as shown in Figure 160.

Figure 160. EFC Manager icon

By double-clicking on this icon, we get the login window for the EFC Manager. We have already defined the user on the EFC Manager on the EFC server. Now we use our username and password and the IP address of the EFC server to login, as shown in Figure 161.

Figure 161. EFC Manager login window

After a successful login, we can move on to managing the fabric.

7.3 Managing the ED-5000 fabric

In our environment, we now have an ED-5000 and the EFC server installed; both have valid IP addresses and are accessible through the corporate intranet. Now, we are going to use the newly created user ID and the EFC

Manager software on the remote workstation to perform these tasks as listed in the following sections:

- 7.3.1, “Identifying the ED-5000 to the EFC Manager” on page 204
- 7.3.3, “Configuring the ED-5000 with the Product Manager” on page 208
- 7.3.4, “Configuring the FC ports” on page 211
- 7.3.5, “Configuring the FC operating parameters” on page 214

7.3.1 Identifying the ED-5000 to the EFC Manager

After logging on to the EFC server, the Product View opens with no devices installed, as shown in Figure 162.

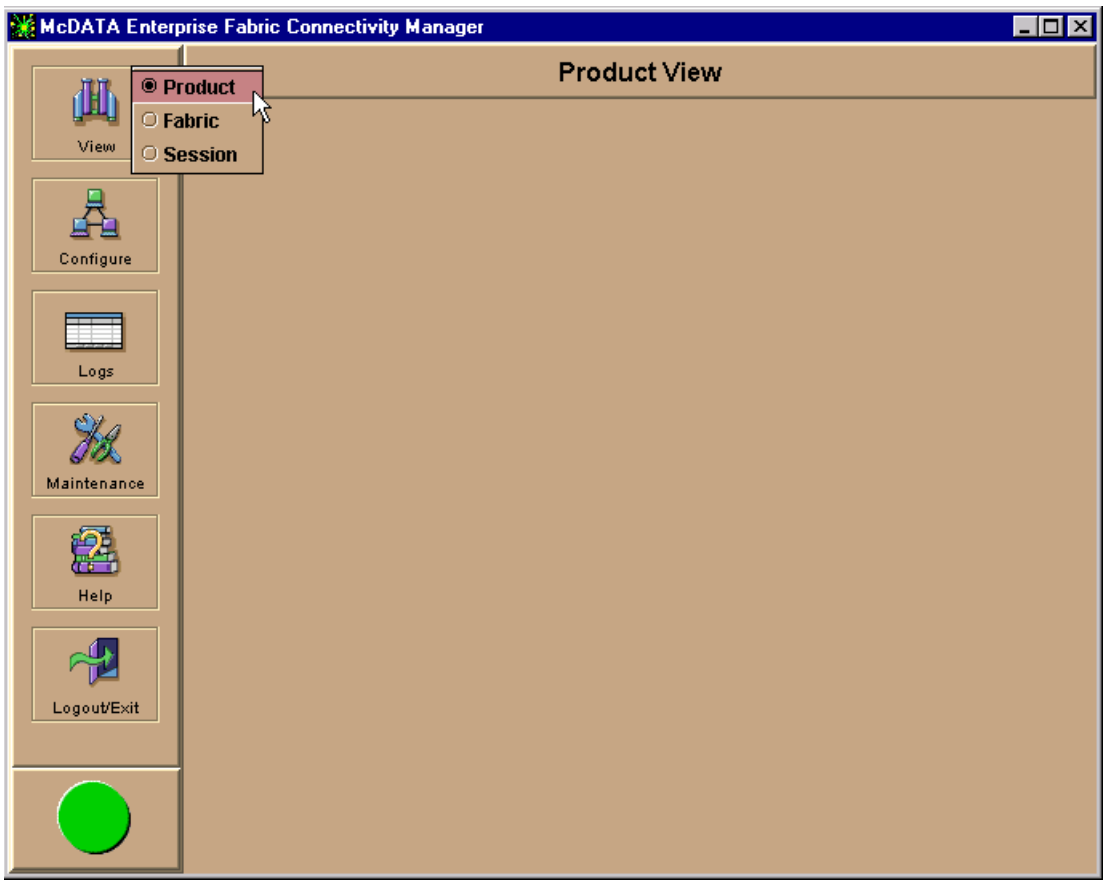

Figure 162. Product View with no devices installed

To the left, there is a panel from where we can access the different management and monitoring functions. To identify the ED-5000 to the EFC

Manager to be managed, we need to tell the EFC Manager the IP address of the ED-5000. This is accomplished by selecting **Configure -> New product**, as shown in Figure 163.

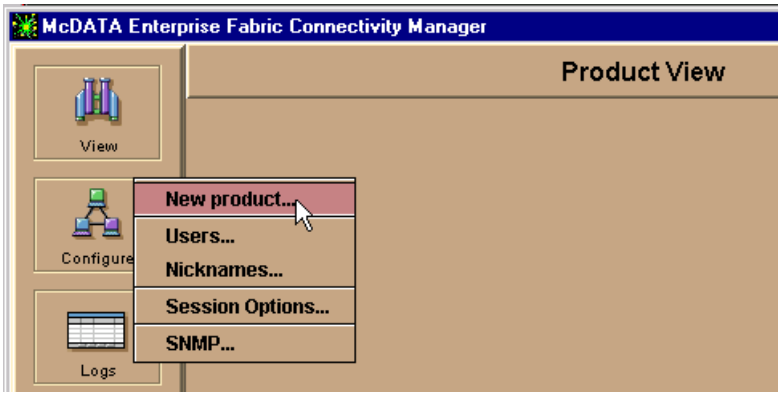

Figure 163. Adding new product

Selecting this takes us to the New Product entry field, where we have to fill in the IP address of the director that we want to add. This is shown in Figure 164.

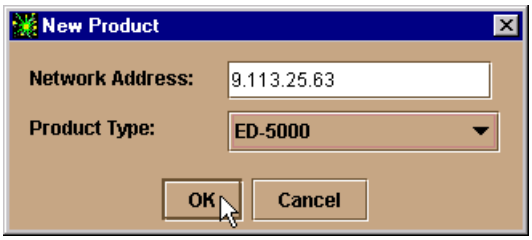

Figure 164. Adding new ED-5000 with its IP address

The ED-5000 was properly installed in the network previously, and now the EFC server can communicate with it. Therefore, the new director appears as an icon in the left area of the main window, as shown in Figure 165. We can also see the IP address of the ED-5000. A green circle indicates that the switch is up and running with no known problems.

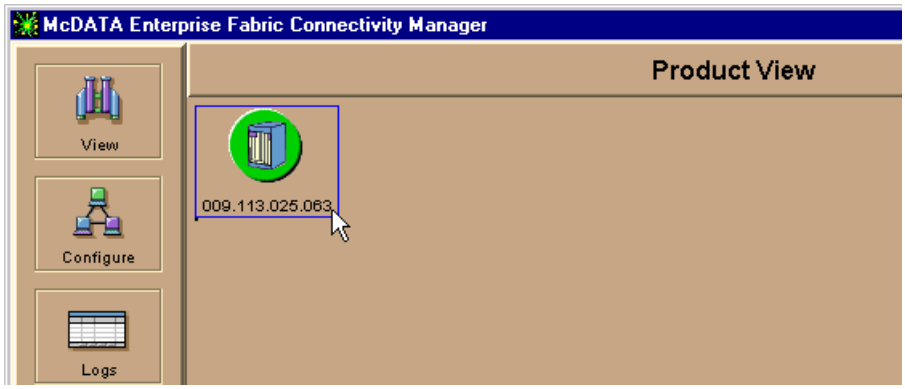

Figure 165. New ED-5000 icon

7.3.2 Assigning nicknames to WWN

As with IP addresses and the DNS, life can be made easier by defining nicknames for WWN. The names can be the DNS host name, in the case of only one adapter in one host. If there are more adapters in one device, we recommend that the nickname consists of the hostname and some extension to distinguish between them. This helps later when we have to identify devices, for instance, when we setup zoning.

This is especially true because the ESS has Emulex adapters built in. We have other Emulex adapters in our SAN, so it would be useful to distinguish between the ones in workstations and the ones found in the ESS. For the hosts we use the nickname that we used in the ESS configuration. For the ESS we chose, as an example RPKA93_C1B2A1. This means the ESS RPKA93 with the FC port in Cluster one, Bay two, and Adapter one.

We use the information to include specific adapters in a zone. This does not mean that we are able to restrict access to volumes on the ESS by doing this. On the ESS, every volume is accessible through every FC port and other steps have to be carried out on the ESS itself to accomplish this. We detail this in Chapter 4, “Configuring the ESS with native Fibre Channel” on page 63. Including ESS FC adapters within specific zones might be useful when we want to influence the bandwidth that a specific group of devices (zone) can get, and through which bay we want the data to go.

For our example, we configure a nickname for one RS6000 and two for the ESS. This is done by selecting **Configure -> Nicknames**, which opens the window without any nicknames configured. We use the Add button to add some nicknames, as shown in Figure 166.

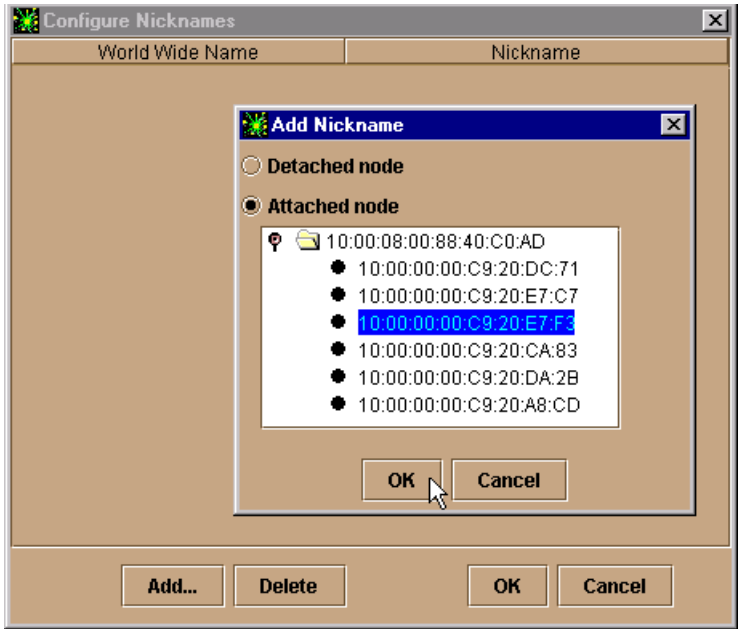

Figure 166. Configuring nicknames

After adding some nicknames, the window looks like Figure 167. Doing this is similar to a 'hosts' file on a TCP/IP host. The next step to a more user friendly naming of the WWNs can be a dedicated service similar to the DNS.

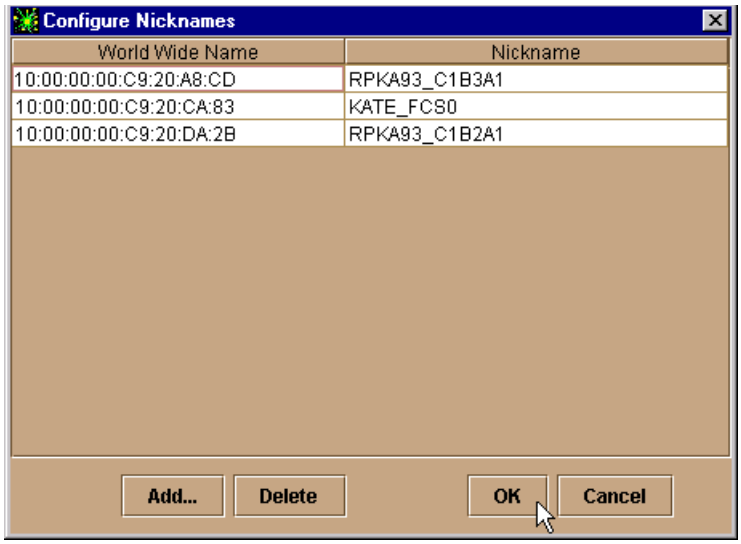

Figure 167. Nickname window with nicknames assigned

In our simple case, it might be easier to work with the WWN and to skip the task of assigning nicknames. However, as more devices attach, maintaining the fabric with names is more convenient than figuring out which WWN belongs to which machine.

After assigning the nicknames, the Node List View shows the names of those that are currently attached, as shown in Figure 168. There are not many devices attached, but with a growing SAN it becomes more and more important to distinguish between the nodes.

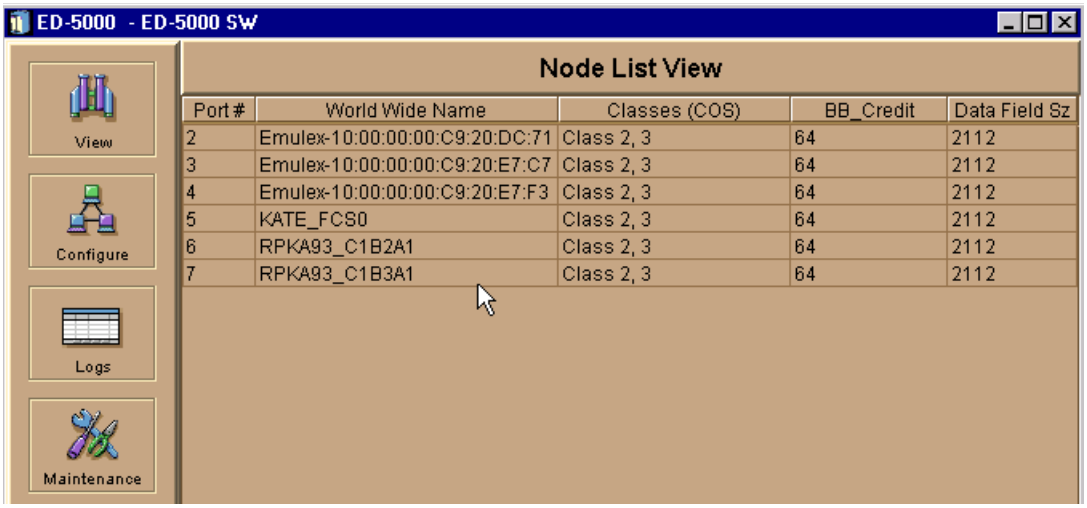

Figure 168. Node List View with some nicknames

7.3.3 Configuring the ED-5000 with the Product Manager

Now, we can configure the ED-5000. We click on the ED-5000 icon in the Product View. This opens the Product Manager in its own window with the Hardware View, as shown in Figure 169.

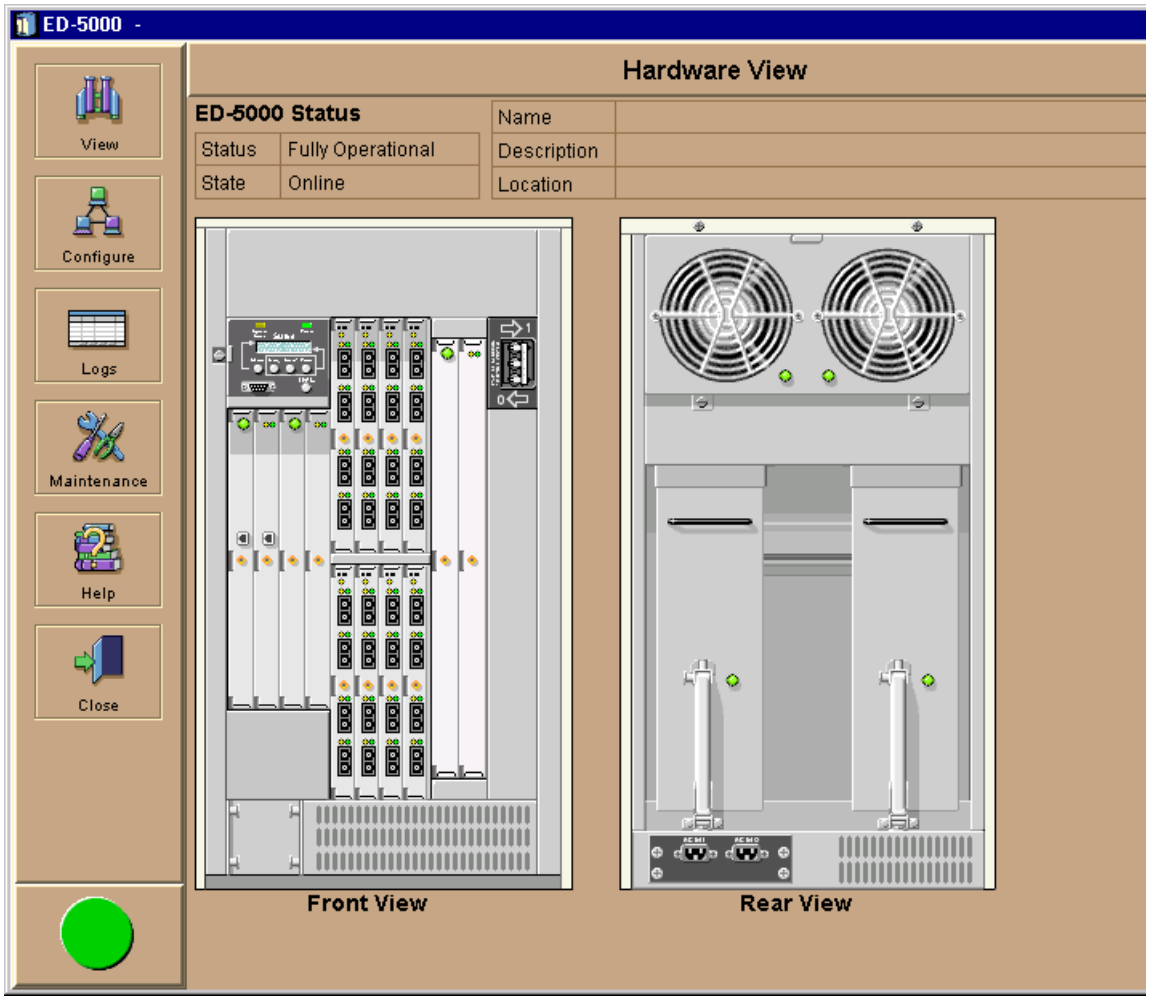

Figure 169. Hardware view of the ED-5000

On the main window, we see the operational status of the switch itself. The switch is fully operational and online. Also, there are fields for the switch name, the switch description, and the location in the main window. This is useful to distinguish among a number of installed directors, if installed.

To configure this information we select **Configure -> Identification**, and we are presented with a dialog window with data entry fields, as shown in Figure 170.

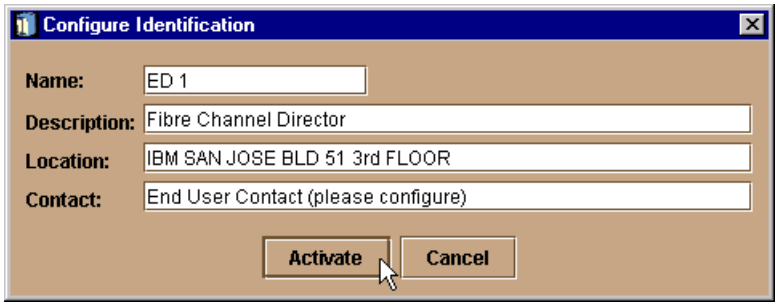

Figure 170. Setting the switch identification

After activation, the display of the main window changes and places the name of the switch in the title bar, and the name, description and location displayed in the window, as shown in Figure 171. This information is used in various locations of the Product Manager to identify the selected director. The name of the ED-5000 can be the same as the host name assigned to the DNS. This prevents too many names for the same director.

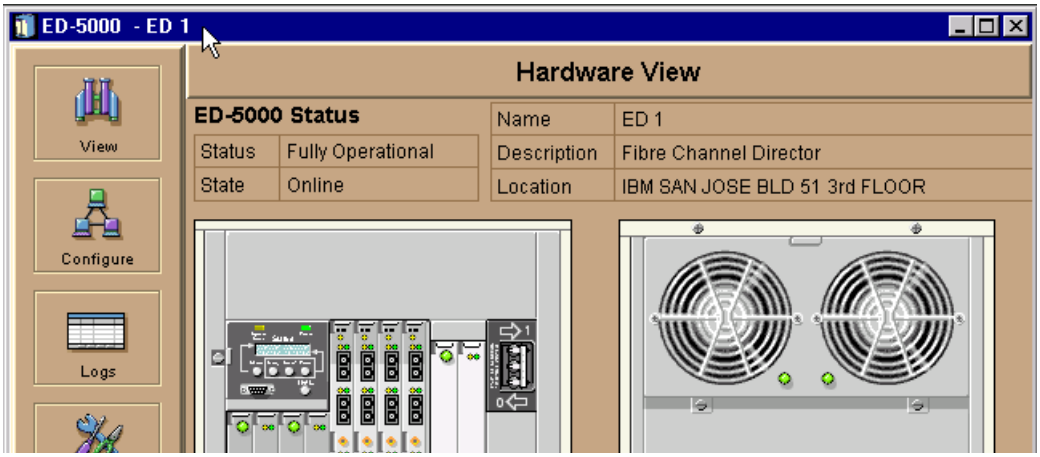

Figure 171. Switch information

By selecting the Hardware View, we are presented with the front and the back of the selected director. These are interactive views, where you can query the status of the FRUs by clicking on them. Once clicked, more information of the FRU is displayed with indicators relating to its operational status. For example, clicking on one of the G_Port modules opens the Port Card View. Now, we can select a port by clicking on it and get detailed port information similar to that shown in Figure 172.

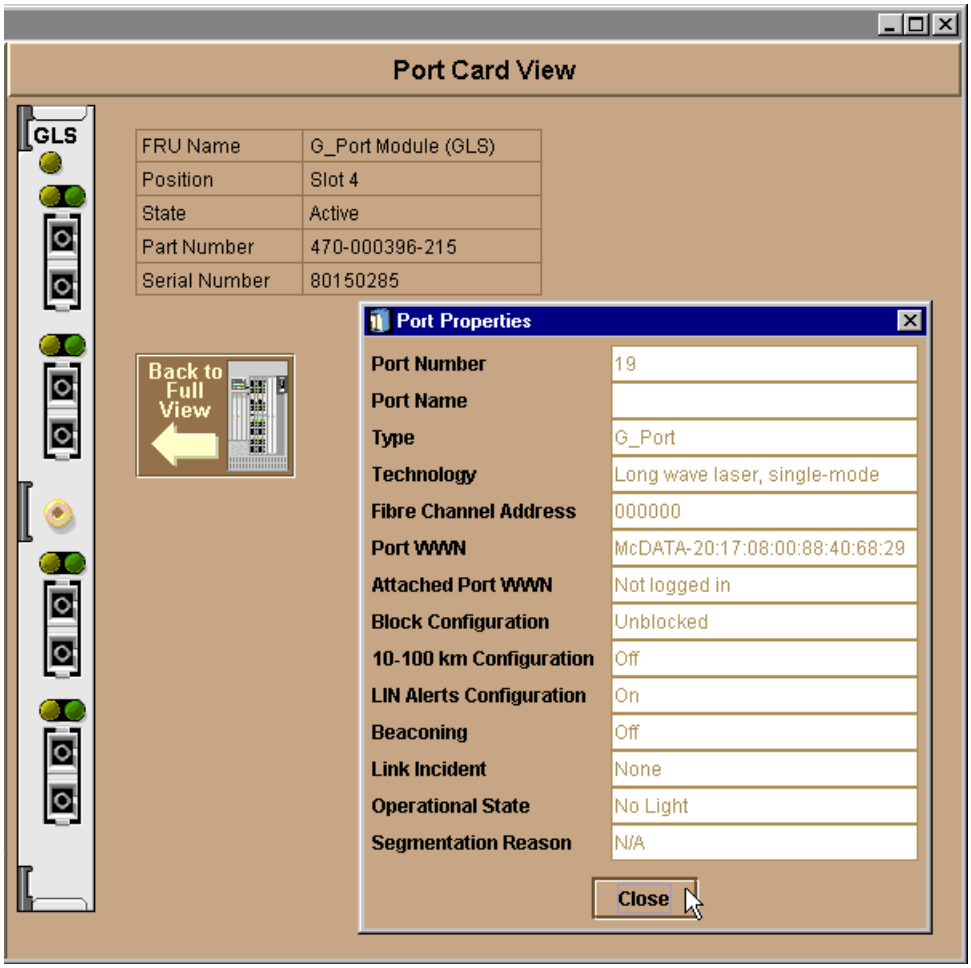

Figure 172. Port Card View with Port Properties

7.3.4 Configuring the FC ports

To configure the options relating to each port, we select **Configure -> Ports**. We are now presented with the Configuration Ports window, which is shown in Figure 173.

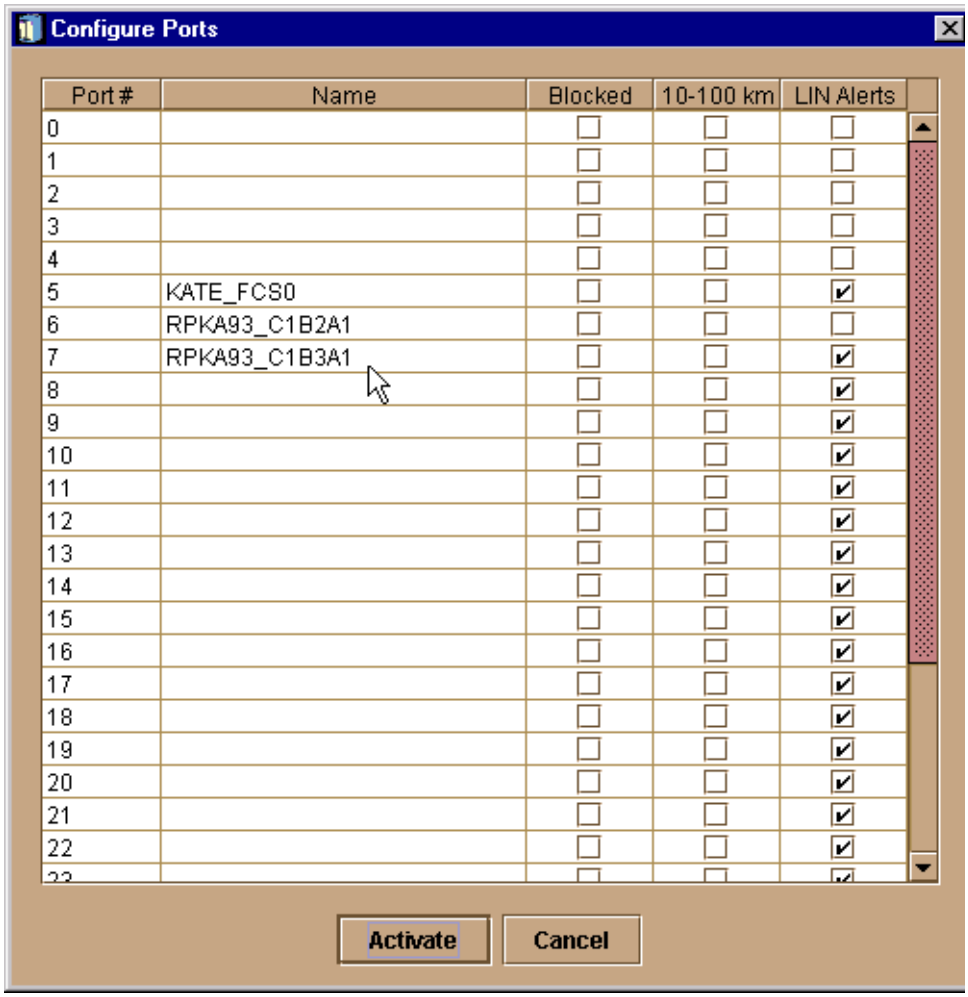

Figure 173. Configure Ports window

The port number is automatically assigned and cannot be changed. We can specify a port name here, but this is only useful if the cabling on the port does not change that often. We recommend using the name used as the WWN nickname. To save the information use the Activate button.

Of more interest here is the ability to block specific ports, to use extended distance buffering, and disable link incident (LIN) alerts. A link incident is a problem on a link, which is visible in the Link Incident Log. It is indicated with a little yellow triangle next to the port, as shown in Figure 174. To view the LIN log, go to **Logs -> Link Incident Log**, as shown in Figure 174.

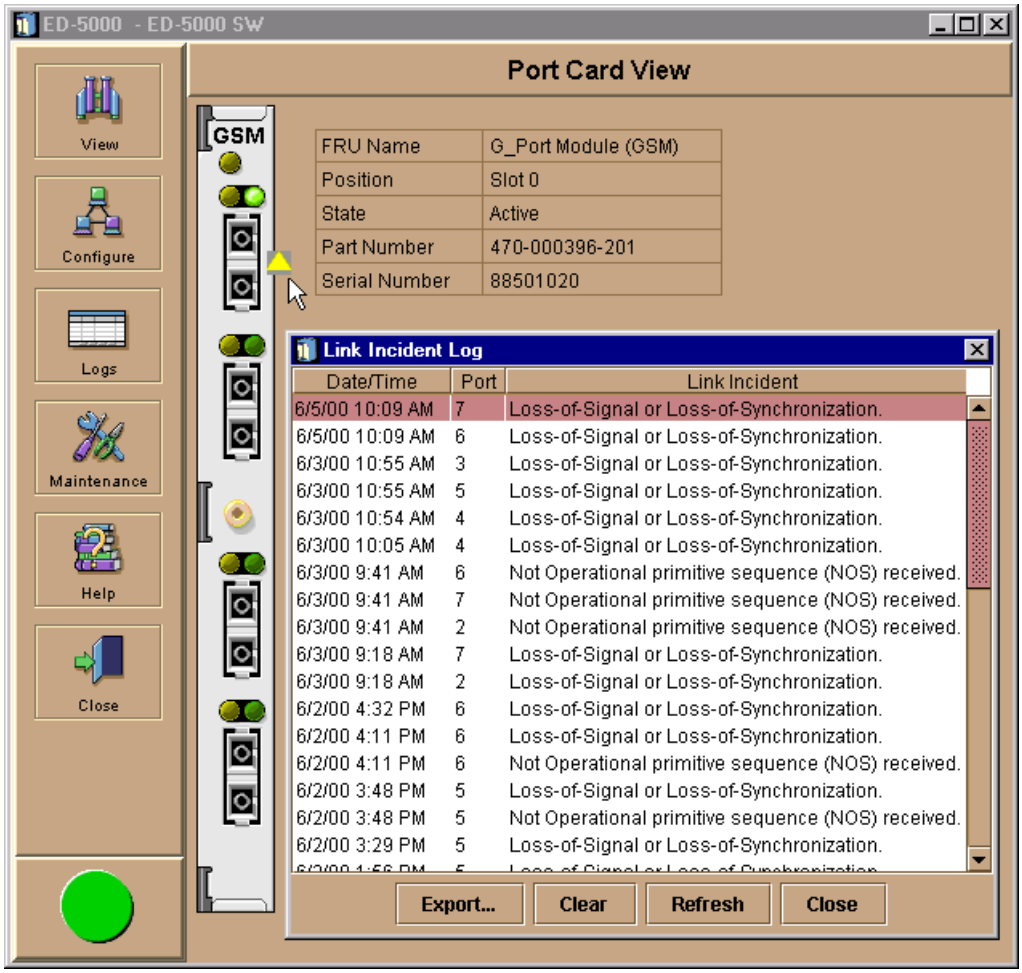

Figure 174. Link Incident Log

To see the changes that have been made to the ports, we select the Port List View, as shown in Figure 175.

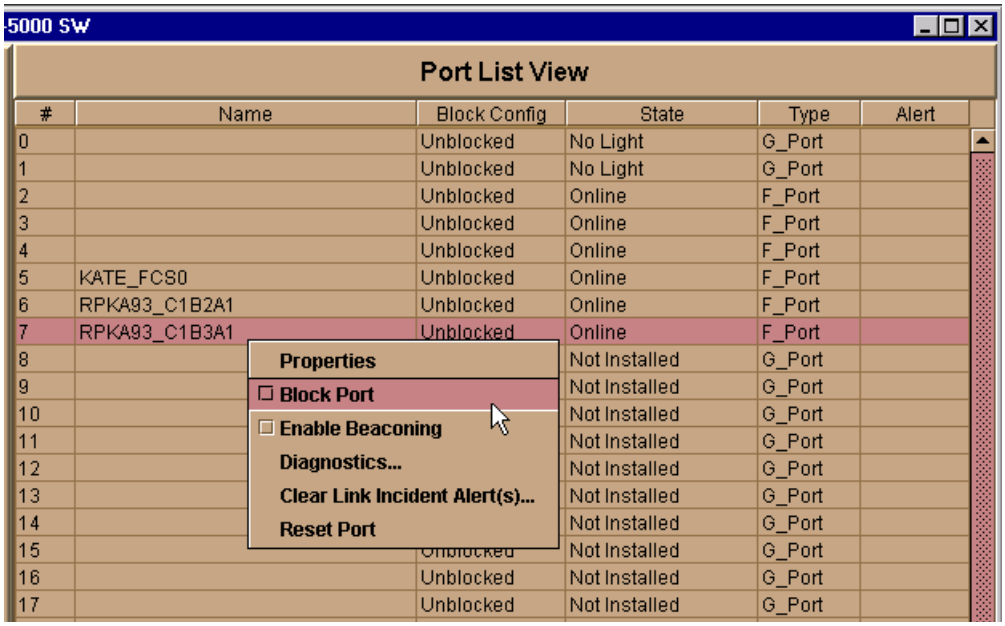

Figure 175. Port list view

7.3.5 Configuring the FC operating parameters

The Fibre Channel operating parameters do not normally have to be changed. The most important setting here is the Preferred Domain ID. This has to be unique for each director within a multiswitch fabric. If we add directors to our SAN, we have to change this parameter. Otherwise, the fabric build process will fail.

To change the operating parameters, we first have to set the ED-5000 offline. To set the director offline, which will terminate all FC operations, we select **Maintenance -> Set Online State**, which is shown in Figure 176.

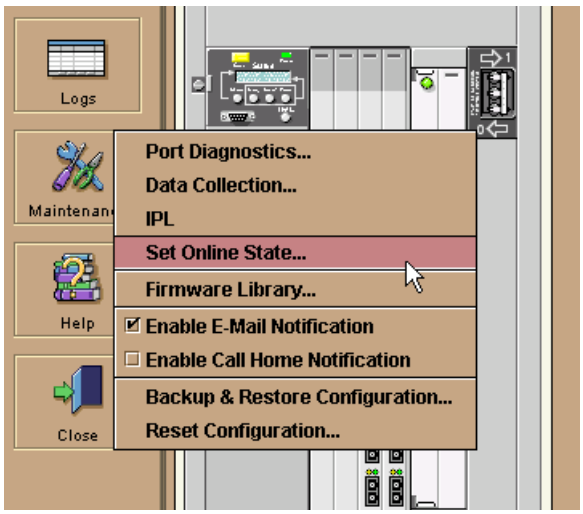

Figure 176. Set Online State

Now, we can go to the Configure Operating Parameters window by selecting **Configure -> Operating Parameters**. Here, we can change some of the Fibre Channel parameters for the director, for example, the preferred domain ID. This is shown in Figure 177.

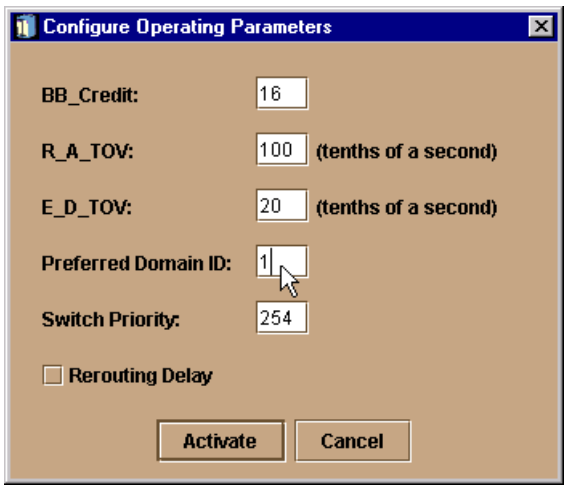

Figure 177. Configure operating parameters

For more information regarding the other management and monitoring functions of the ED-5000, refer to the *Enterprise Fabric Connectivity Manager User Manual.*

Now the director is ready for use in the Fibre Channel network. It can be connected to devices, such as other switches, storage, or hosts.

7.4 Zoning with McDATA

The advantage of a SAN, and the accessibility of any storage anywhere, raises issues which must be solved. For example, do we really want any device to have the ability to access all storage? Zoning helps to split the SAN into logical parts.

The following sections cover these topics:

7.4.1 Zoning overview

Zoning is a feature that is used to divide a fabric into groups of devices. Members of a zone can only communicate with members in the same zone. For instance, zoning is used to separate groups of hosts with different operating systems from each other, or to separate part of the fabric for testing, or for separating user groups for security and management reasons.

There are two approaches to assign devices to a zone. One is commonly called hard zoning and uses the FC ports and the devices connected to it to build a zone. The other is called soft zoning and uses the WWN of the HBA of a host or the FC interface adapter of another FC device to build the zone.

The two approaches to configure devices for a zone can be mixed.

WWN based zoning has the advantage that we can rearrange the devices on the ports without affecting the definition of the zones. On the other hand, we have to change the zoning definition every time we change a member in a zone, or if we have to replace an HBA or an interface adapter to a storage device. Port based zoning requires us to change the zoning definition every time we change the wiring to the devices, but we can change devices connected to ports without having to redefine the zones.

7.4.2 Preparing to define zones

To view and manage the zones in the McDATA fabric, we must open the zoning view of the EFC Fabric Manager. To do so we connect to the EFC Manager, then select **View -> Fabric**, as shown in Figure 178.

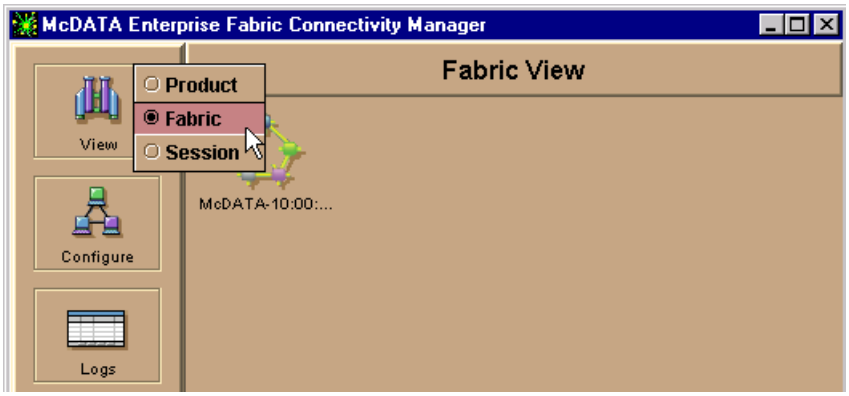

Figure 178. Opening Fabric View of the EFC Manager

Click on the fabric icon in the upper left of the Fabric View to open the EFC Fabric Manager. Now the Fabric Manager opens as a second window with the Topology View. In our case, as there is only one McDATA installed, and there is only one switch visible, as shown in Figure 179.

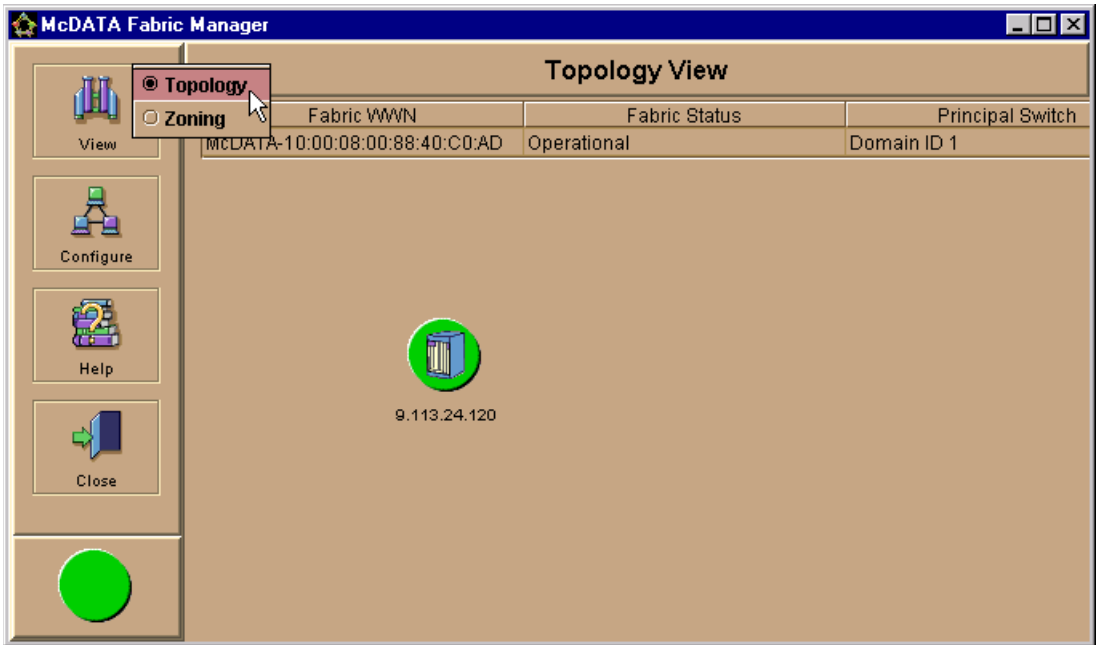

Figure 179. Topology View of the Fabric Manager

To change to the Zoning View, we again select the View button and this is also shown in Figure 179.

In Figure 180, we show the Zoning View of the Fabric Manager. This is the entry point for viewing and managing zones within the selected McDATA fabric. At this moment in time, there are no active zone sets. In the topics that follow we show how to set up zoning in our environment.

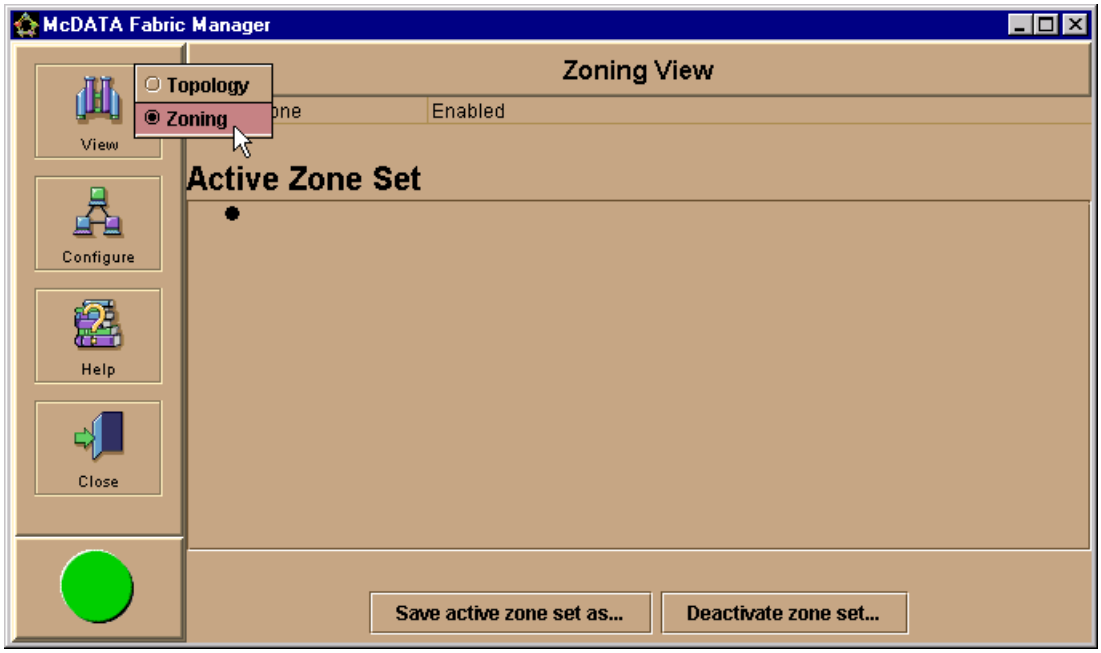

Figure 180. Zoning View of the Fabric Manager

7.4.3 McDATA zoning concept

With the McDATA we can specify zone sets. A zone set is a group of zones that can be activated and deactivated at the same time. This can be used to save different configurations for different tasks, for example, if we want to have different devices in the same zone for backup, but not during normal operation. Devices that are not configured in a zone within the active zone set are considered as members of the default zone. Also, if no zone sets are activated, all devices are in the default zone. With the default zone enabled, it is possible for all devices in the default zone to communicate with each other in parallel to the currently active zone set. You can also disable the default zone independently from the active zone.

An example of how zones and zone sets are related is shown in Figure 181.

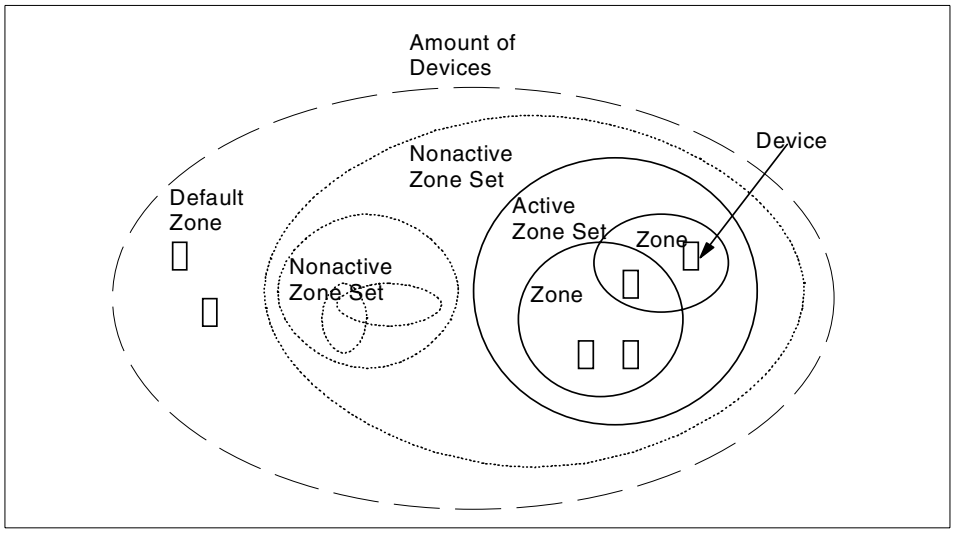

Figure 181. Example for McDATA Zoning

7.4.4 Creating a zone set and establishing an NT zone

The next step in our setup of the McDATA SAN is shown in Figure 182. There are hosts already connected to it, but we introduce the zoning concept — first by establishing an NT zone, and then later by the addition of an AIX zone.

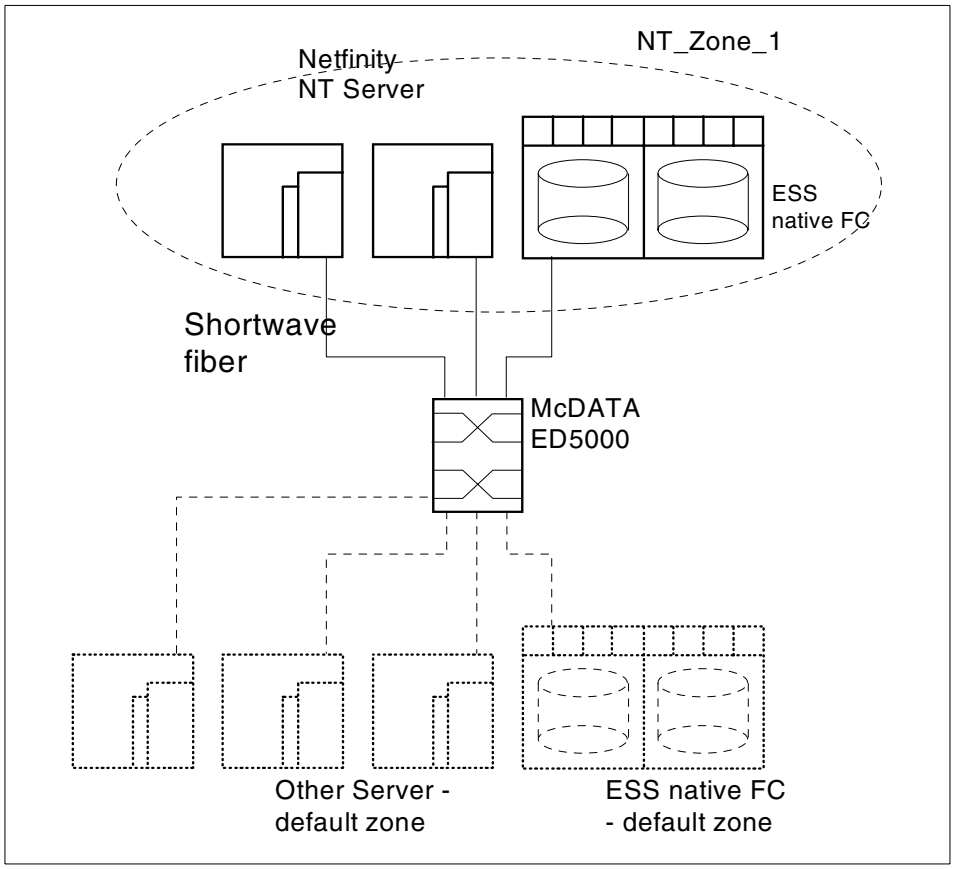

Figure 182. NT zone with McDATA

To create a new zone set, select **Configure -> Zone Sets**. This displays the zone set library and provides us with options for changing our zone definitions, for example, creating new zone sets, deleting zone sets, or modifying existing zone sets as shown in Figure 183.

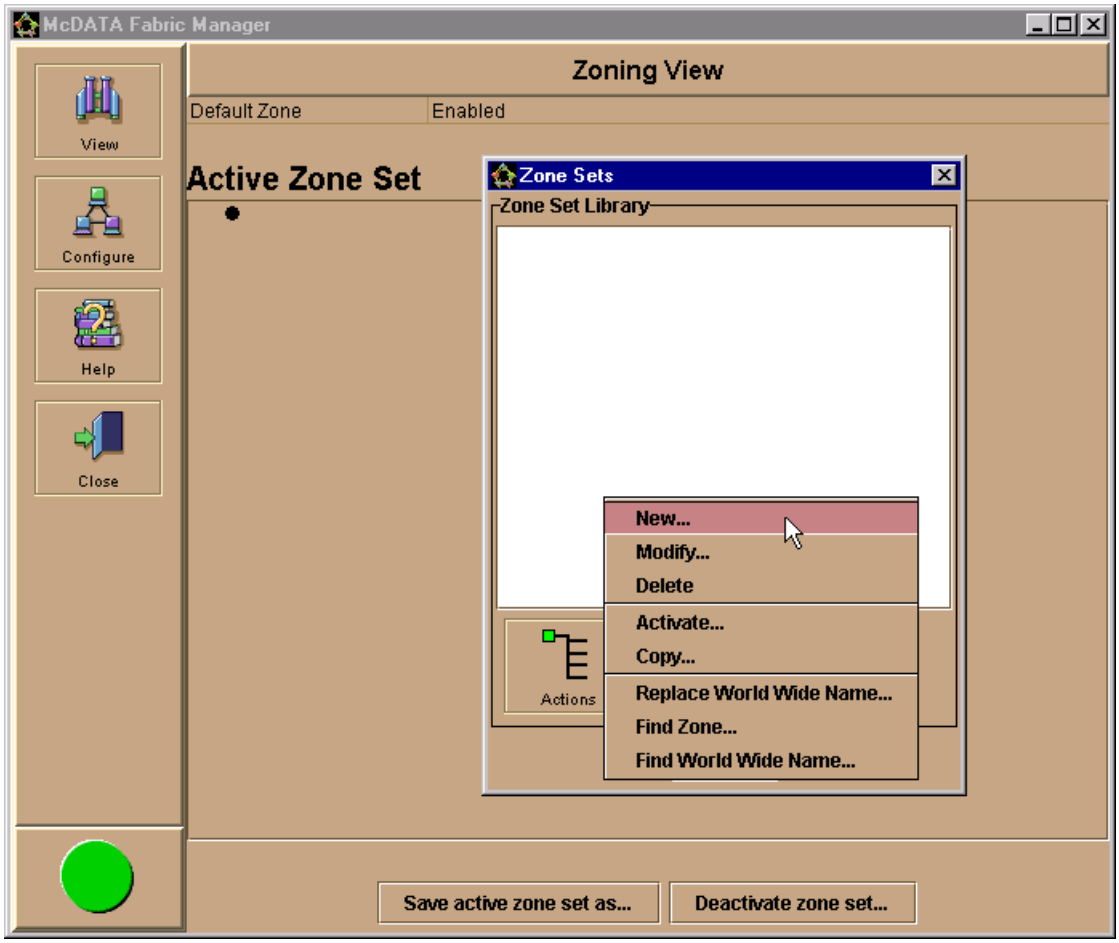

Figure 183. Actions for Zone Sets

Because there are no zone sets in the zone library, we will create one. Starting from the **Actions** button in the **Zone Sets** window, we select **New**. We are presented with a window to define a zone set. From this window we will assign a zone set name. The **Actions** button provides us with the different tasks to maintain the zones in this zone set, as shown in Figure 184. Because there is no zone defined in the zone set, we can only assign a new zone to the zone set.

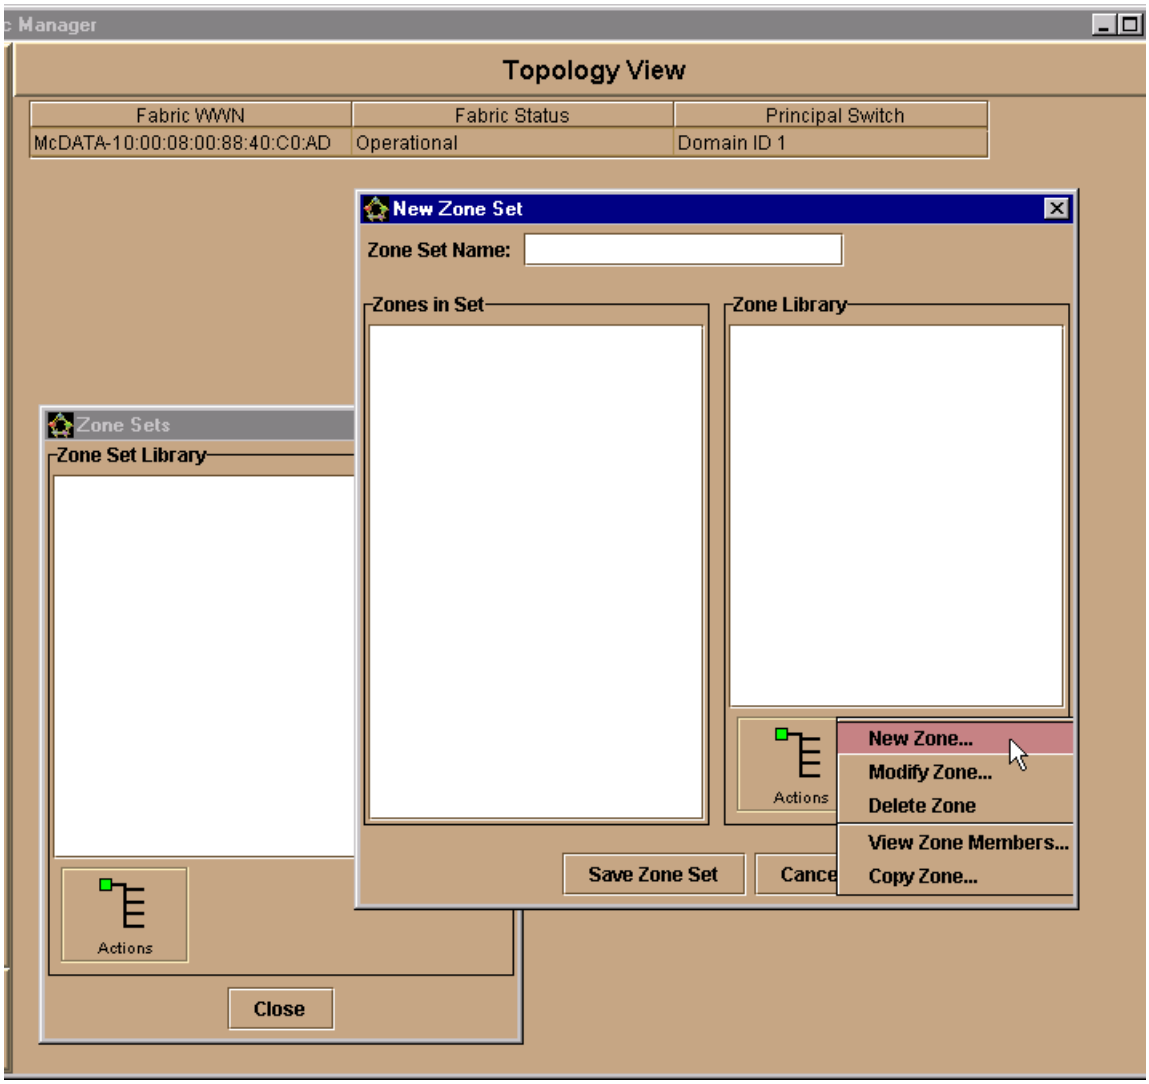

Figure 184. Defining a new zone set

Now a window opens where we can define the members of the zone, as shown in Figure 185. We want to define a zone that includes our Windows NT machines and one ESS. The ESS is already connected to the switch and the hosts are currently detached.

To the right of the window in Figure 185, we are presented with all of the McDATA connected Fibre Channel adapters and their WWNs. To the left we can view all the members currently defined for the zone. There are two check

boxes where we can choose if we want to assign a McDATA port to a zone, or if we want to specify directly the machines based on the WWN of their Fibre Channel adapter. In our example, we chose to add the Windows NT servers, regardless of which port they are connected to.

To assign a Fibre Channel adapter to the zone, we drag and drop the WWN associated with the Fibre Channel adapter of our Windows NT server to the left part of the window. This is what we will do with the ESS.

The NT host we want to define for this zone is currently detached, but we can also define this host using this window. To perform this we use the **Actions** button and select **Add Detached Node**. We get a data entry field where we insert the WWN of the node as shown in Figure 185.

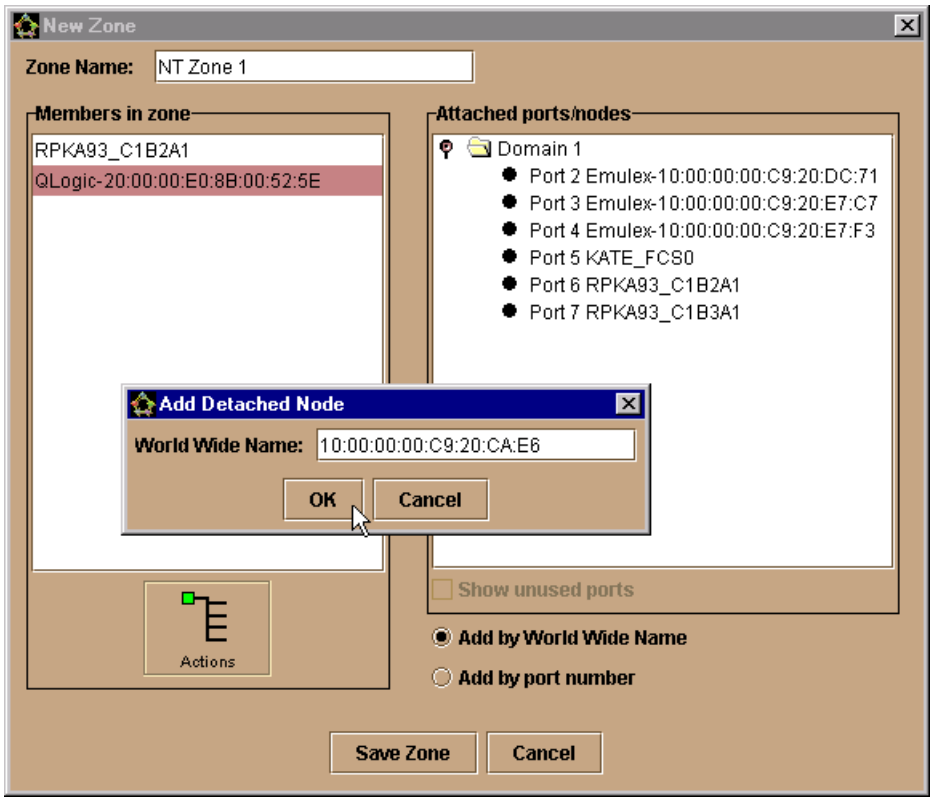

Figure 185. Adding a detached Fibre Channel node

While trying to save the zone using the **Save Zone** button, we got an error message. This is because of the spaces in the zone name, as shown in

Figure 186. This is not allowed, so to circumvent this problem we changed it to NT_Zone_1.

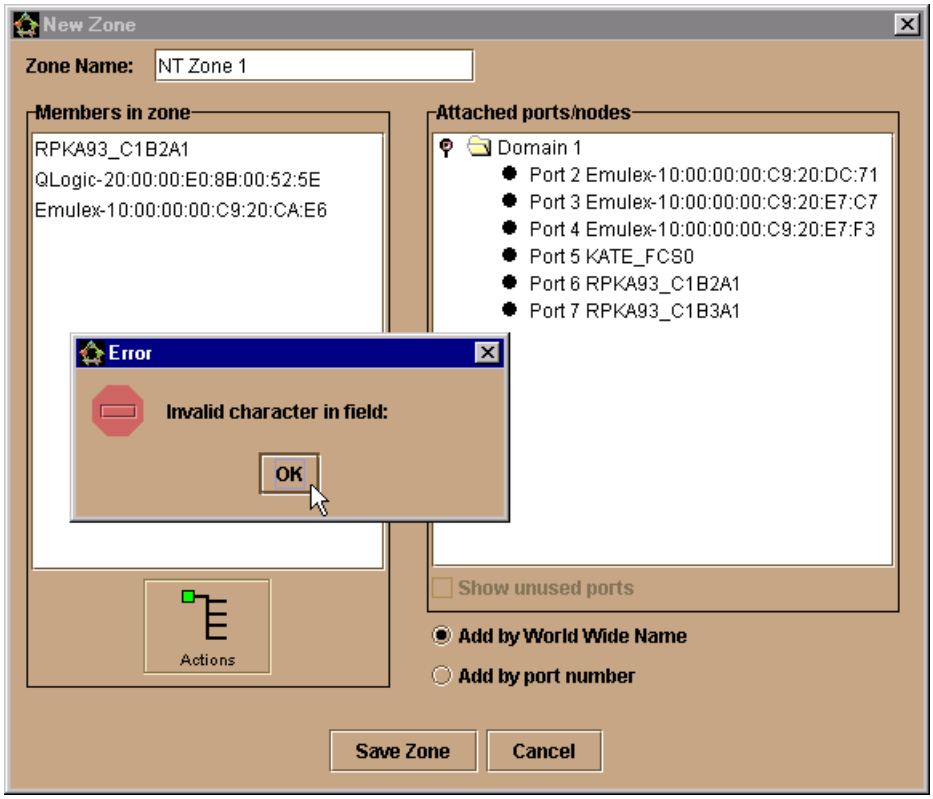

Figure 186. Defining the zone

We illustrate that we can specify the members in one zone based on the WWN, and other members based on the port they are connected to, as shown in Figure 187.

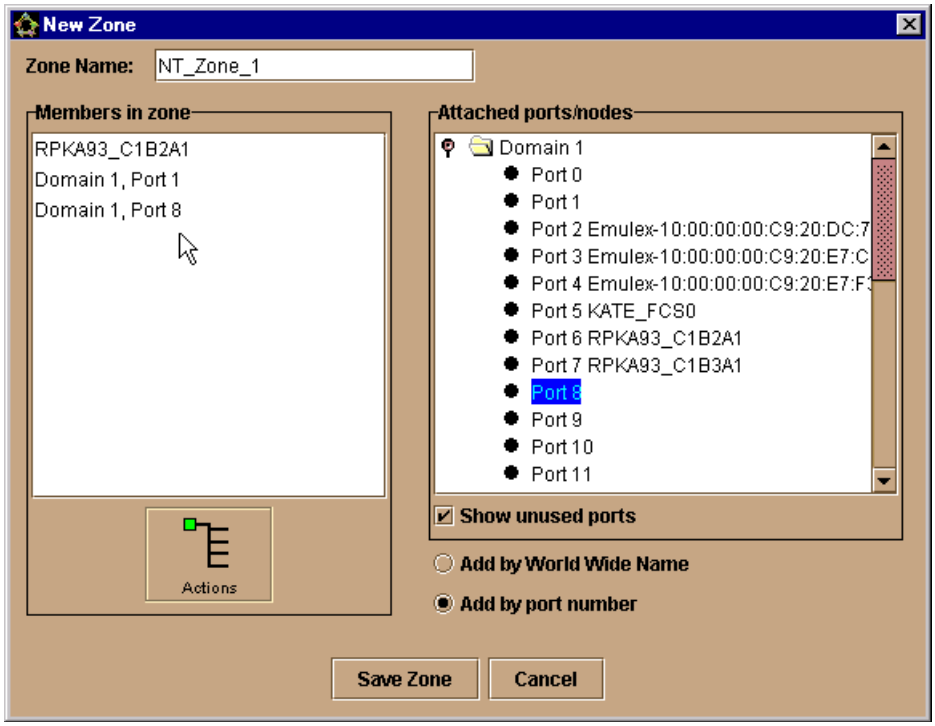

Figure 187. Members in zone based on WWN and port

After selecting **Save Zone**, we return to the New Zone Set window where we can view the zones in the Zone Library in the right half of the window.

To view the zone members, highlight the zone in the Zone Library and then go to **Actions -> View Zone Members**, as shown in Figure 188.

Figure 188. Viewing zone members after creating a zone

To assign the newly created zone, NT_Zone_1, in the zone set we created before, we drag and drop the zone from the Zone Library to the Zones in Set on the left side in the window. This displays the zones assigned for this set, as shown in Figure 189.

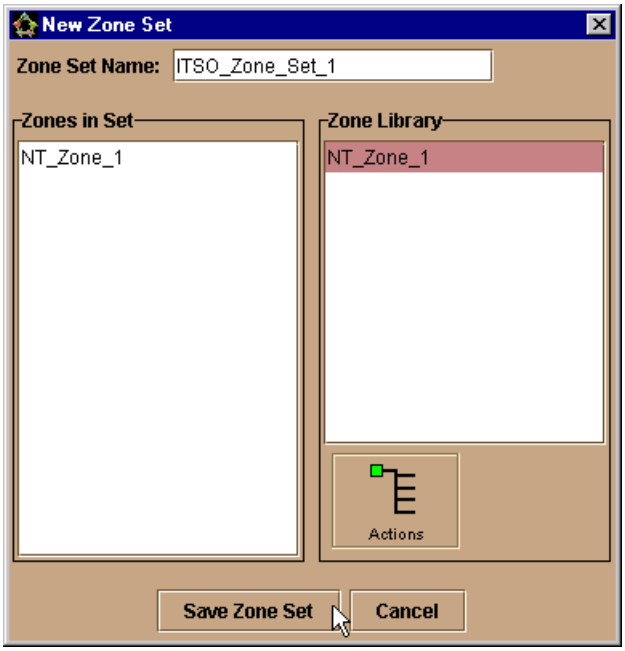

Figure 189. Assigning a zone to a zone set

Now we save the zone set as a member in the Zone Set Library with the **Save Zone Set** button. The Zone Set Library now looks like Figure 190. We have one zone set, ITSO_Zone_Set_1, with one zone, NT_Zone_1. We have one ESS port and two hosts within the zone, one host attached with an Emulex adapter and one host detached with a QLogic adapter.

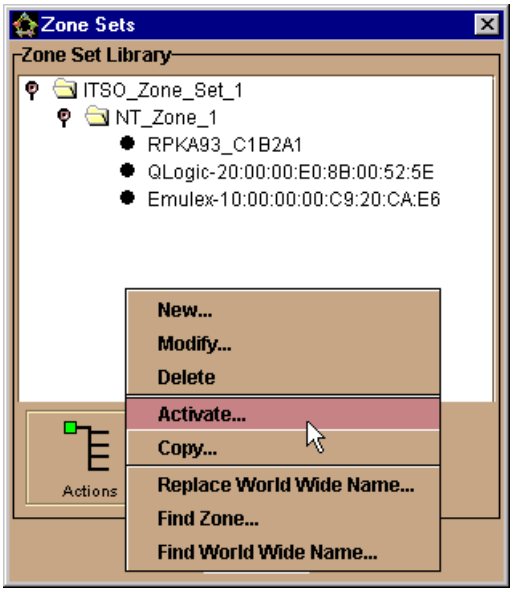

Figure 190. Zone Set Library with one zone set and one zone with two hosts

The ITSO_Zone_Set_1 zone set, with NT_Zone_1 zone, is now defined. However, it is not active. We can define more zones in the zone set or create other zone sets. To finish our zoning example, we will activate the zone set now. This is also done using the **Action** button and by clicking **Activate**. With this action, we are prompted to start or to cancel the activation of the zone set with the Activate Zone Set window. We **Start** it and receive a message that the activation is complete, as shown in Figure 191.

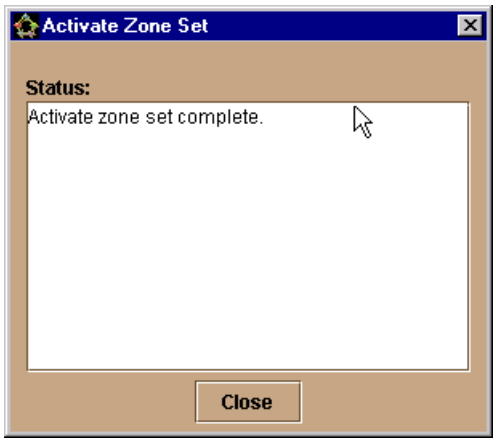

Figure 191. Activate Zone Set complete

After returning to the Zoning View window and expanding the zone set by clicking the small symbol to the left of ITSO_Zone_Set_1, and then NT_Zone_1, the Zoning View of our fabric manager looks like Figure 192.

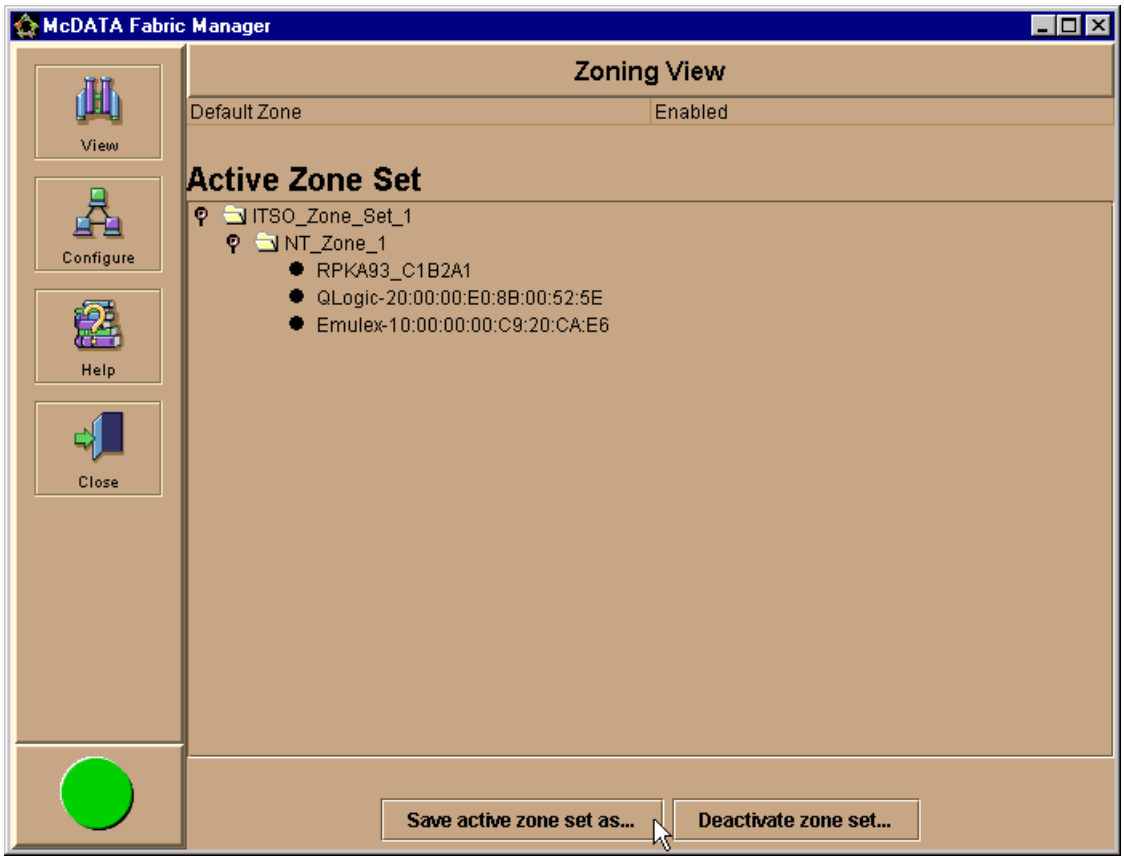

Figure 192. Active Zone Set shown in the Zoning View of the fabric manager

We see the zone set and the associated zone by its name and the configured host adapters by their manufacturer's name and their WWN. Also we see the nickname for the ESS port which we configured earlier. If we had not configured a nickname for the ESS port, we would only see it as another Emulex adapter, which are used as the Fibre Channel adapters in the ESS.

From within the Zoning View window, we can also manipulate the zone sets, for example, deactivating a zone set or saving the zone set. As an example, we copy the same zone set, but assign it a different name, as shown in Figure 193.

Figure 193. Saving zone set with different name

7.4.5 Adding an AIX zone to the existing zone set

We also have AIX hosts and we want to add a zone with the AIX systems and another ESS, as shown in Figure 194. The AIX hosts are already connected to the McDATA, so all we need to do is to define the zone in the zone set which already includes the NT zone.

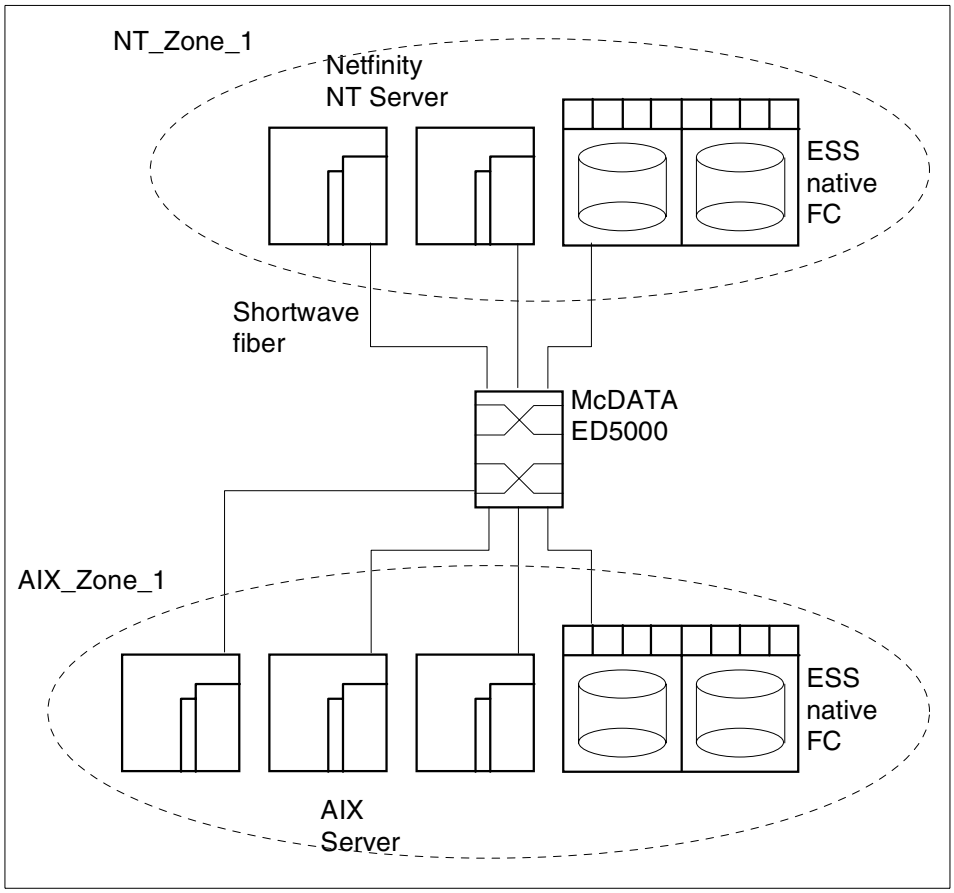

Figure 194. NT and AIX zones with McDATA

To add another zone to the zone set, we basically follow the steps described previously. We navigate to the Zone Set Library and mark the zone set in which we want the zone to be added. In our case this is the 'ITSO_Zone_Set_1'. Then we use **Modify** from the **Actions** menu, which allows us to change the content of the zone set, as shown in Figure 195.

.

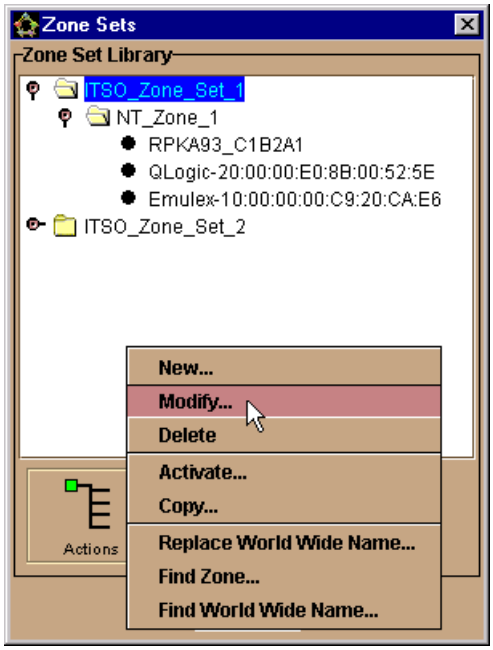

Figure 195. Modify zone set

From the Modify Zone Set window we select **Actions -> NewZone**, which opens the New Zone window. We drag and drop the adapters of our AIX hosts to the Members in zone entry field, as shown in Figure 196. The last entry to the left of the window is another ESS, but which is without a nickname.

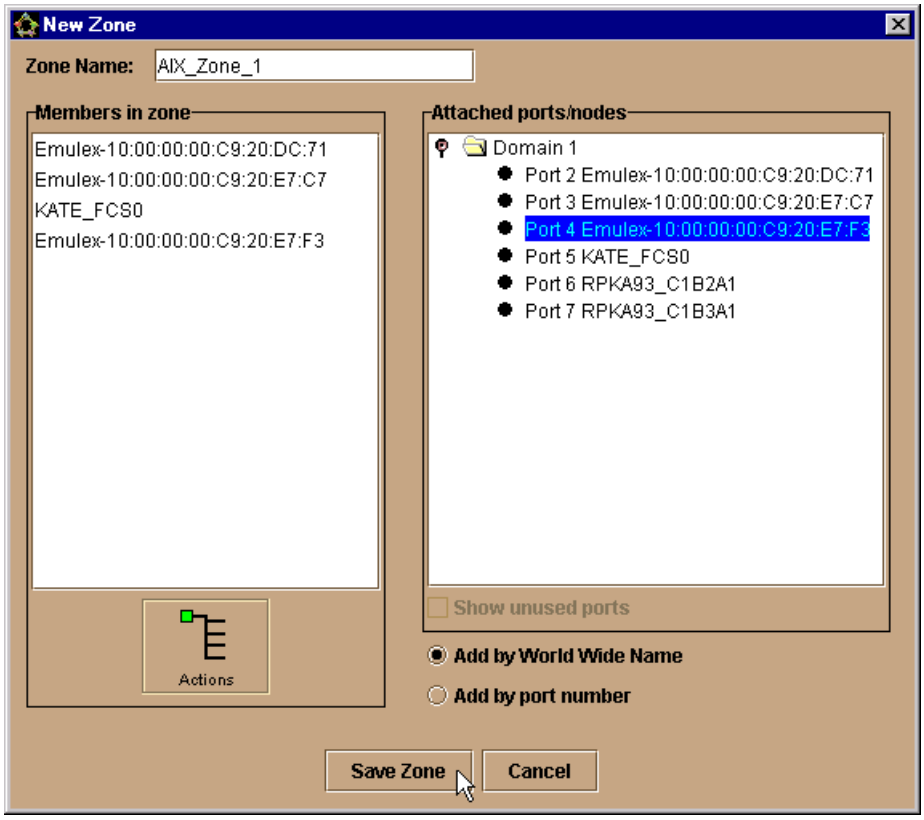

Figure 196. Defining an AIX zone

After selecting **Save Zone**, we assign this zone to our zone set. Just as we did previously with the NT zone, we drag and drop the AIX zone in the Modify Zone Set window from the Zone Library to the Zones in Set, on the left side in the window, as shown in Figure 197.

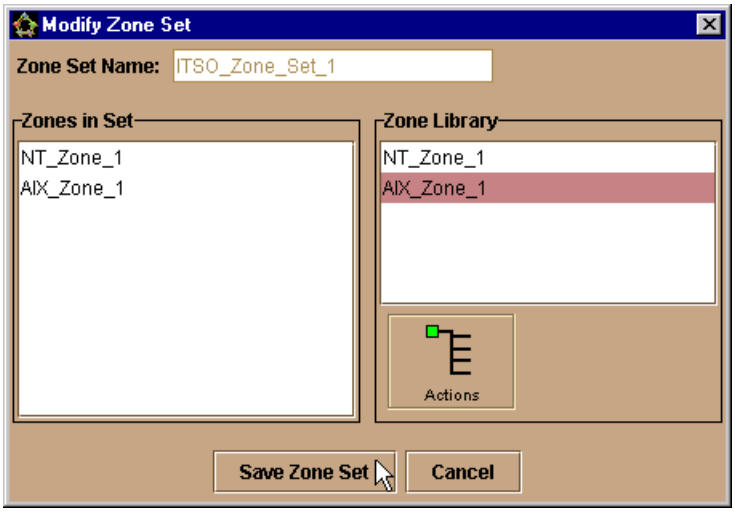

Figure 197. Assigning the AIX zone to the zone set

After saving the zone set, we have two zones, the NT_Zone_1 and the AIX_Zone_1. Both of these are in our ITSO_Zone_Set_1 zone set, as shown in Figure 198. To apply the changes in the zone set, we must activate it by selecting **Actions -> Activate**, as shown in Figure 198.

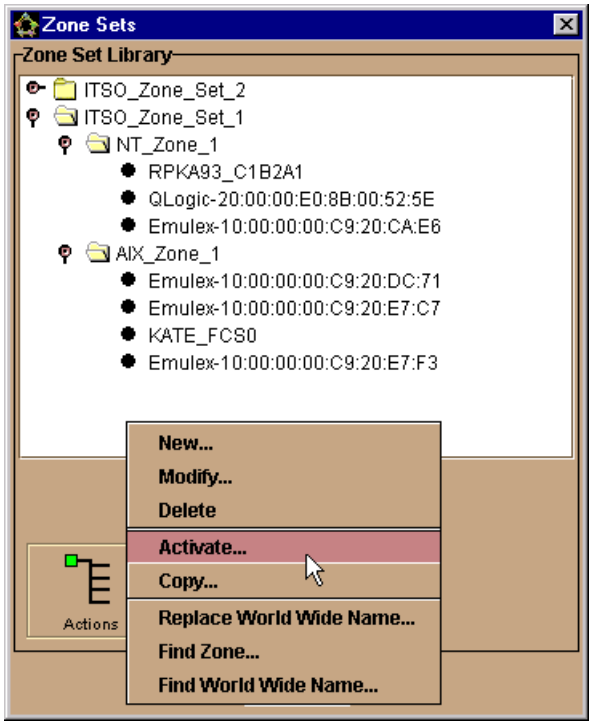

Figure 198. Zone Sets with two zone sets and two zones in one set

By activating or deactivating this zone set, we activate or deactivate both zones at the same time. The result of doing this, after confirming that we want to start the activation and selecting OK on the resulting message, means that we now have a Zoning View, with an Active Zone Set, which looks like Figure 198.

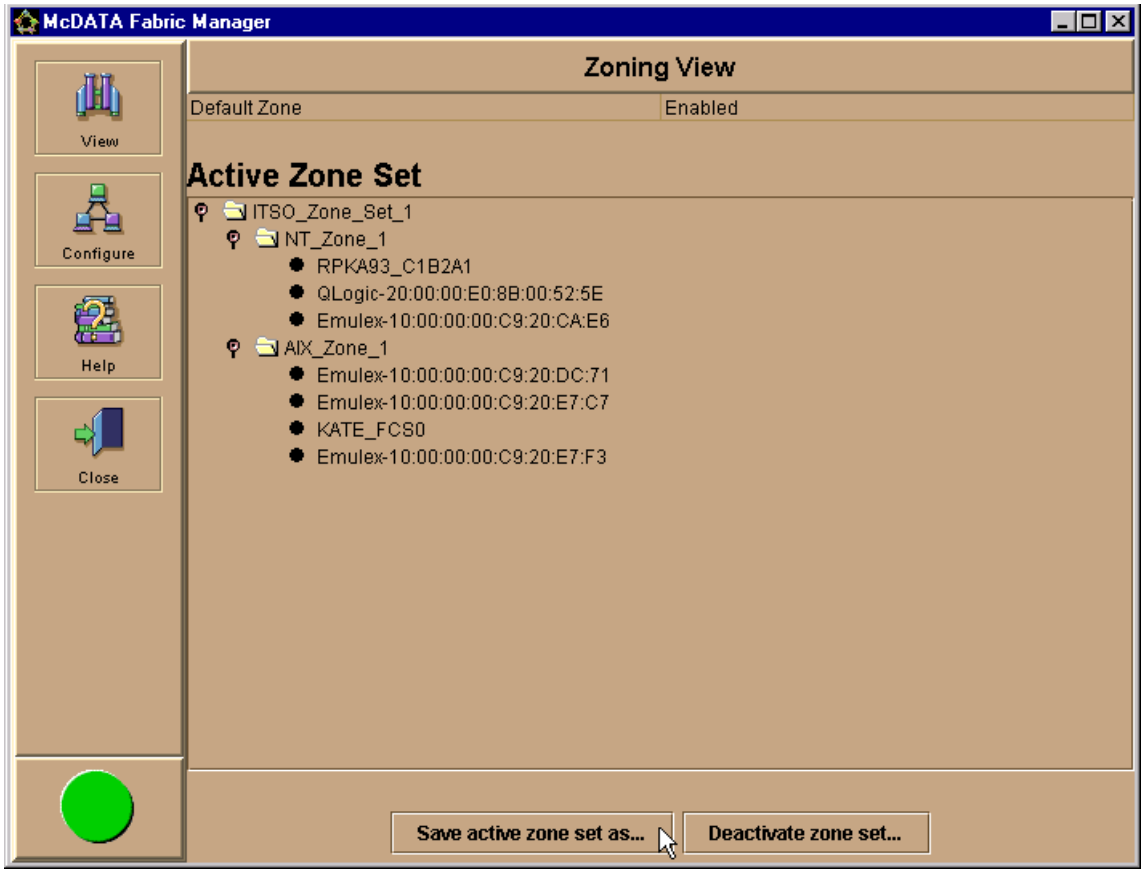

Figure 199. Zoning View with active zone set containing two zones

We successfully created and activated a zone set with two zones.

7.5 Cascading with McDATA - building a multiswitch fabric

The focus of McDATA was, primarily, high availability and reliability, not connectivity and extended fabrics. However, times change and with the introduction of E_Port capability of the McDATA G_Ports, it is now possible to cascade McDATA switches.

In the following sections, we cover these topics:

- 7.5.1, “Multiswitch fabric considerations” on page 238
- 7.5.2, “Setting up our McDATA multiswitch fabric” on page 242

7.5.1 Multiswitch fabric considerations

The planning of multiswitch fabrics depends on many things. Are you going to have a local SAN in one site with up to 32 devices connected? Then you may not want to consider cascading switches. If you want to build a SAN between two sites that are far apart, cascading becomes valued. Also, if you need more devices connected, or if you are looking to introduce extra redundancy, cascading is the only way to achieve this.

Nevertheless, we still might think about whether or not, or to what extent, we want to cascade switches. The reason for this is that by using E_Ports we will sacrifice F_Ports. Also, with an extended fabric, the ISLs can possibly become a bottleneck. This will lead to the use of more ISLs, which means even fewer F_Ports.

What seems easy in the first instance, can get more complicated once we add the zoning concept, load balancing, and any bandwidth issues that may appear.

7.5.1.1 Where multiswitch fabrics are appropriate

Certainly, there are some possible solutions where a multiswitch fabric is needed. For example, disaster recovery solutions that are using a SAN can be built upon a McDATA SAN, but only when using E_Ports to connect directors between two sites. We need directors at both sites to back up one site completely. Disaster recovery *and* high availability can be established together using a multiswitch fabric, and open system hosts using Logical Volume Manager (LVM) mirroring together with clustering software, such as HACMP for AIX or Veritas Cluster Server. Due to the high availability and the many ports of the McDATA ED-5000, two McDATA ED-5000 may be enough.

7.5.1.2 Solutions for high availability and disaster recovery

An example of a solution that provides high availability with disaster recovery, is shown in Figure 200.

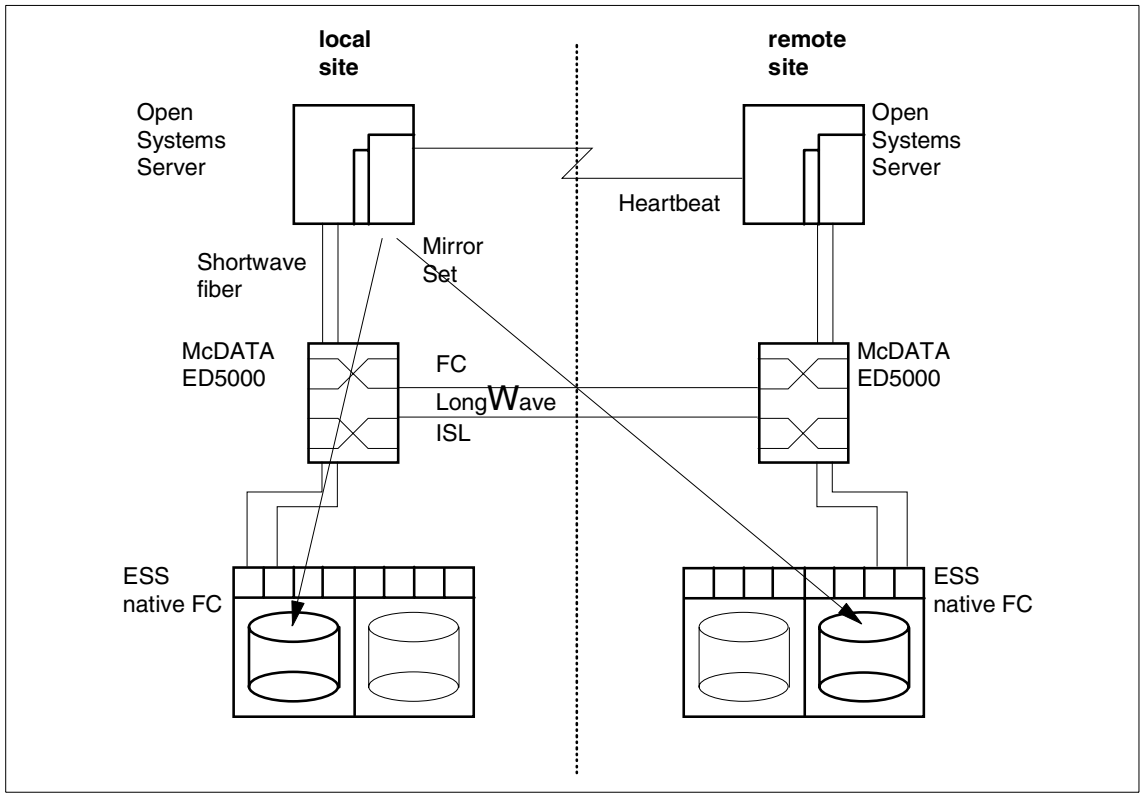

Figure 200. High Availability and disaster recovery with McDATA

This is a setup which consists of the same configuration at both the local and the remote site. Both sites can be up to 10 km apart when using 9 micron fiber-optic cable. The Open Systems Server Cluster, for instance, can consist of two or more RS/6000 with HACMP. The mirroring can be done with the native LVM on AIX.

Another solution can be SUN servers running, for example, the Veritas Cluster Server and the Veritas Volume Manager. Due to the high availability of the McDATA ED-5000, one may be sufficient, if that leaves enough ports to accommodate the specific environment.

When more ports are needed, this solution can easily be extended with another director at each site, which adds even more availability, bandwidth, and ports. This is shown in Figure 201.

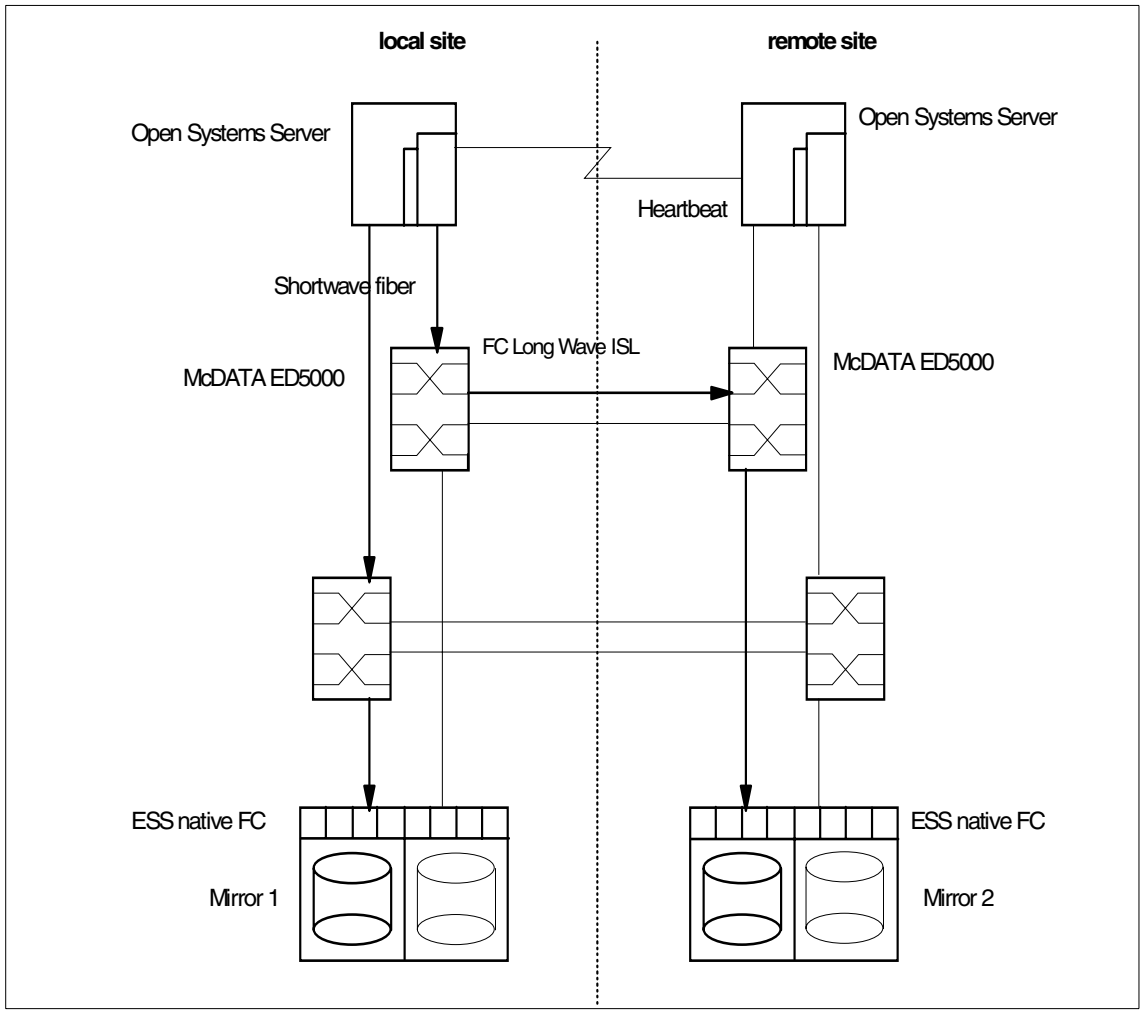

Figure 201. Extended High Availability and disaster recovery with McDATA

The arrows indicate a possible route for the data to get to both parts of the mirrored sets. In this setup there is no single point of failure at a device level, and even if one site completely fails, the other site will take over operation.

In our example for a multiswitch fabric, shown in Figure 202, we are not focusing on clustering. What we want to show is how to apply zoning in a multiswitch fabric.

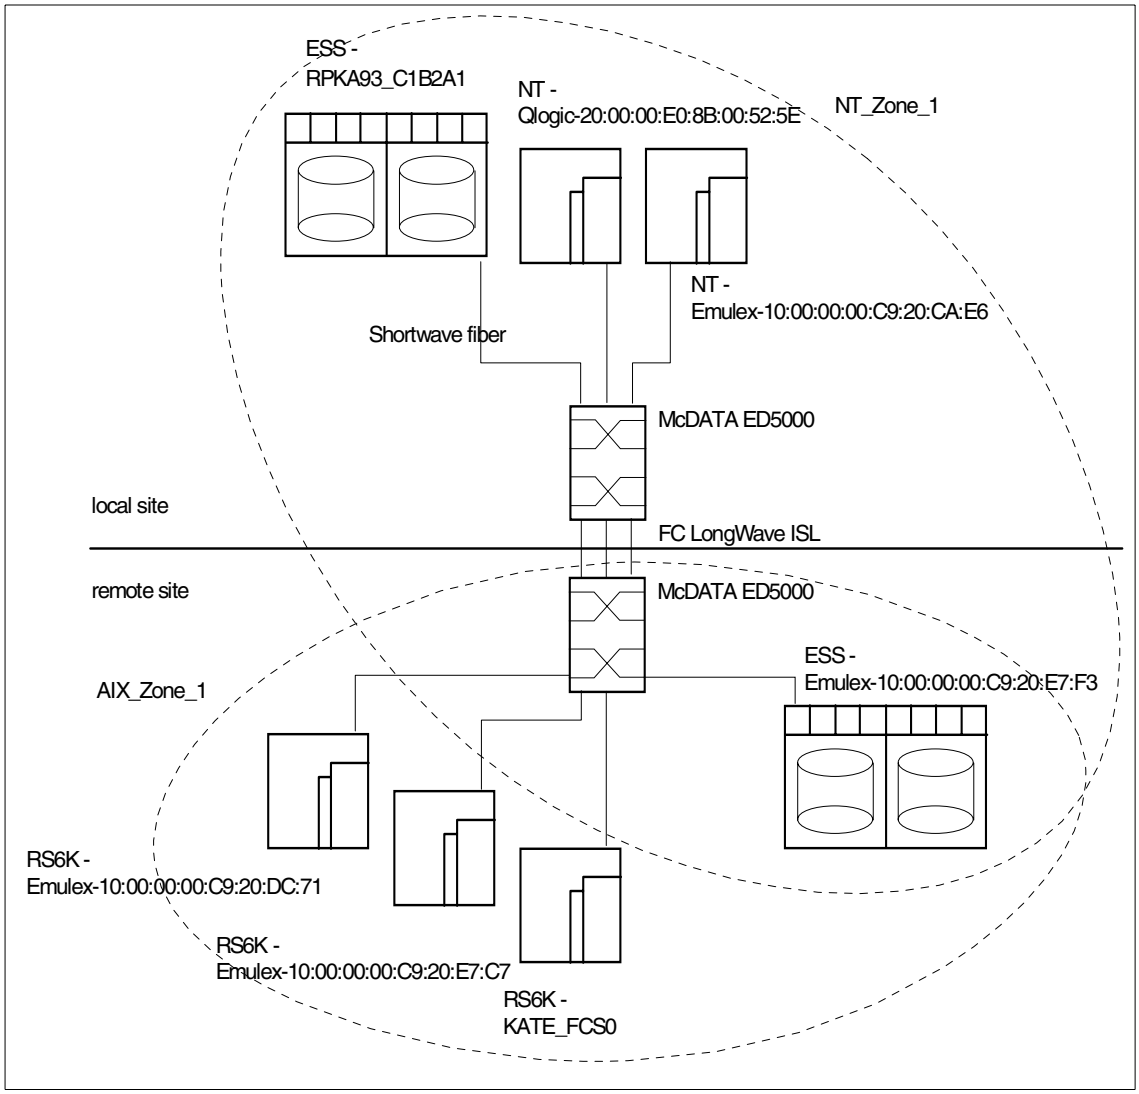

Figure 202. Multi-switch fabric with McDATA

We have our NT zone with two ESSs. One ESS is at the local site and the other is at the remote site. Both sites are connected with three longwave ISLs between the two ED-5000. At the remote site, we have the AIX zone, where the remote ESS is also a member. This example can be used to establish a mirrored set from within the Windows NT Disk Administrator, with one local

copy of the data and one remote. Conversely, the AIX zone is limited to the devices at their site.

7.5.1.3 Limits for the McDATA multiswitch fabrics

McDATA supports only a homogenous environment. This means only McDATA ED-5000 are supported as switching elements in a McDATA SAN. The McDATA fabric supports up to 31 interconnected ED-5000. Although we can connect many directors, the hop count with today's Fibre Channel is limited to two, due to the delay that is applied traversing every switch. The hop count delay increases with every additional switch between the source and the destination.

7.5.2 Setting up our McDATA multiswitch fabric

We will use two ED-5000s for our zoning example. We configure both directors as we did before. First, we define one director with its EFC Server, and then we configure the second director using the same EFC Server.

To include the second switch in the fabric of the first, we basically need to connect the directors with longwave or shortwave Fibre Channel cables. The fabric building process itself is transparent to us as the switches will recognize the connection and automatically configure the G_Ports to be used as E_Ports. However, there are some configuration options that need to be set up, or reviewed before connecting the switches.

7.5.2.1 Setting the domain ID

Each switch itself is recognized in the fabric as a domain and is identified with a domain ID. Every domain ID in the fabric must be unique, ranging from 1 to 31. To view and to change the domain ID, we go to the Configure Operating Parameters window from within the Product Manager of the specific switch. Then we select **Configure -> Operating Parameters**. In the next window, shown in Figure 203, we can change the preferred domain ID and other Fibre Channel parameters for the director, for instance, the Switch Priority.

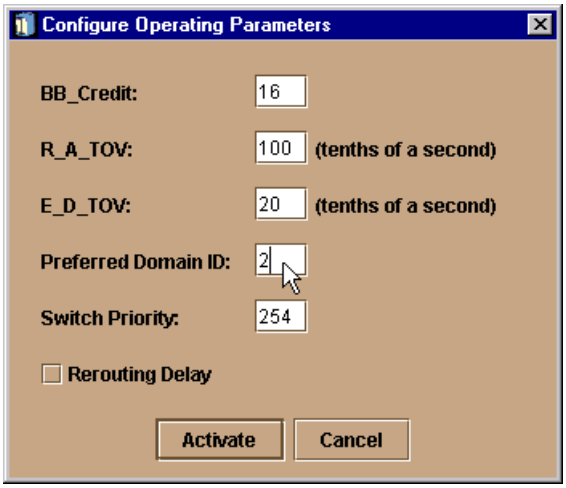

Figure 203. Setting domain ID

7.5.2.2 Setting the switch priority

In every multiswitch fabric, one switch will act with responsibility for the domain address manager functionality. This principal switch controls the allocation and distribution of the domain IDs for all connected switches in the fabric.

The principal switch is the one with the highest switch priority, ranging from 1 (highest priority) to 255 (lowest priority). If switches share the same value for the priority, the one with the lowest number in the WWN becomes the principal switch. To change the switch priority, we also use the Configure Operating Parameters window, shown in Figure 203.

7.5.2.3 Other prerequisites for a multiswitch fabric

To be able to successfully establish a multiswitch fabric some important prerequisites apply. One is the zoning configuration, which must be compatible. What this means is that the active zone set name must be the same, and the zones with the same name must have the same members. Also the operating parameters, resource allocation time out value (R_A_TOV) and error detection time out value (E_D_TOV) must be the same.

7.5.2.4 Configuring the ports for the ISLs

The ports for the ISLs can be configured like the other ports as we described in 7.3.4, “Configuring the FC ports” on page 211. From here we can check the checkbox for extended distance buffering (Figure 204).

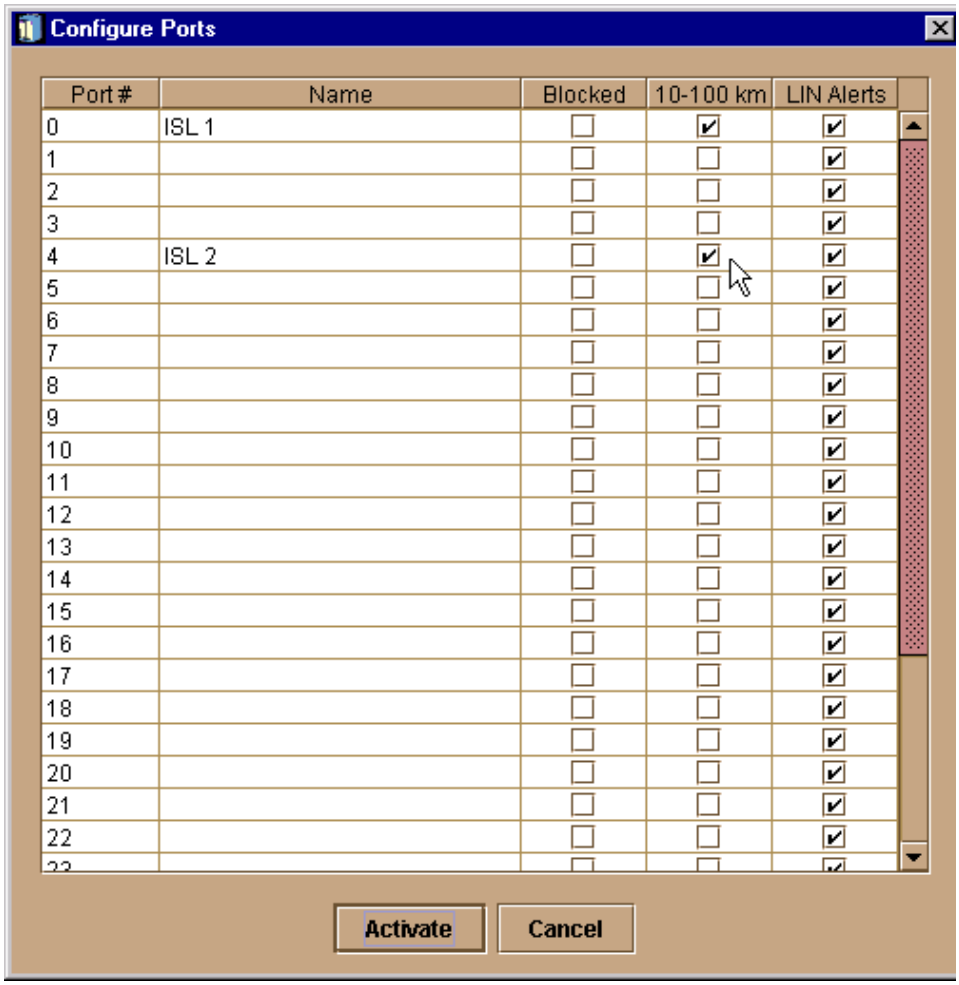

Figure 204. Configure ISL ports

7.5.2.5 Connecting the directors

Now we can connect the two switches with ISLs. We are using two longwave ISLs between the two switches. As a result, we now have two switches in the EFC Manager's Product View, as shown in Figure 205.

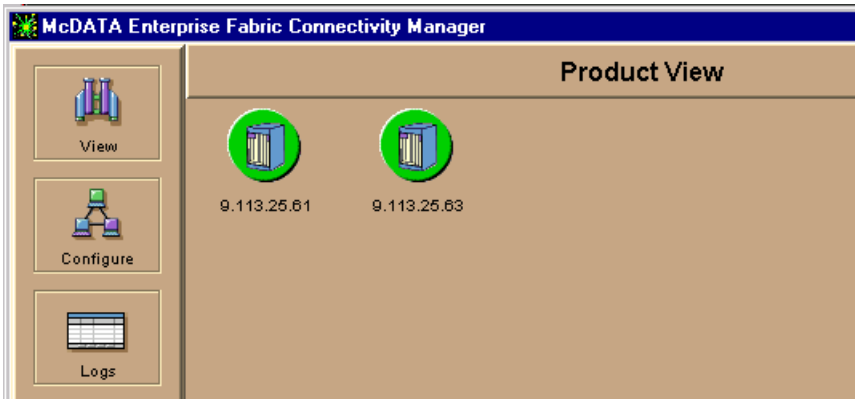

Figure 205. Two managed switches from within one EFC Manager

To see what the topology looks like now, we navigate to the Topology View of the Fabric Manager. Select **View -> Fabric** in the EFC Manager window and click on the icon that represents the fabric in the main window, as shown in Figure 206. The number '2' in the pentagon fabric icon indicates that we have two directors installed in the fabric.

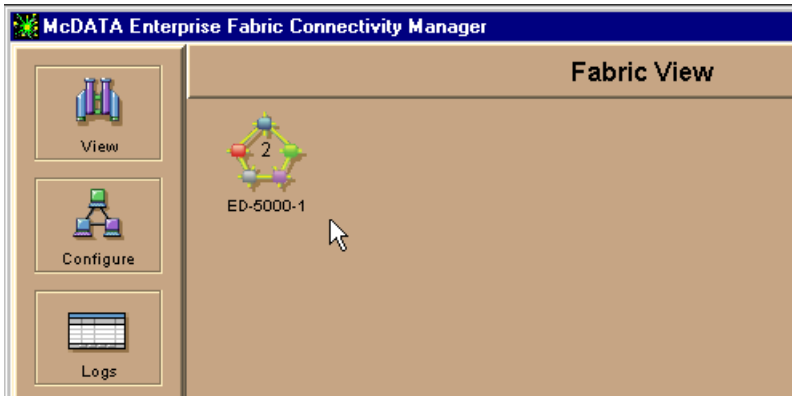

Figure 206. Fabric View

Clicking on this icon will open the Fabric Manager in a dedicated window with the Topology View as shown in Figure 207.

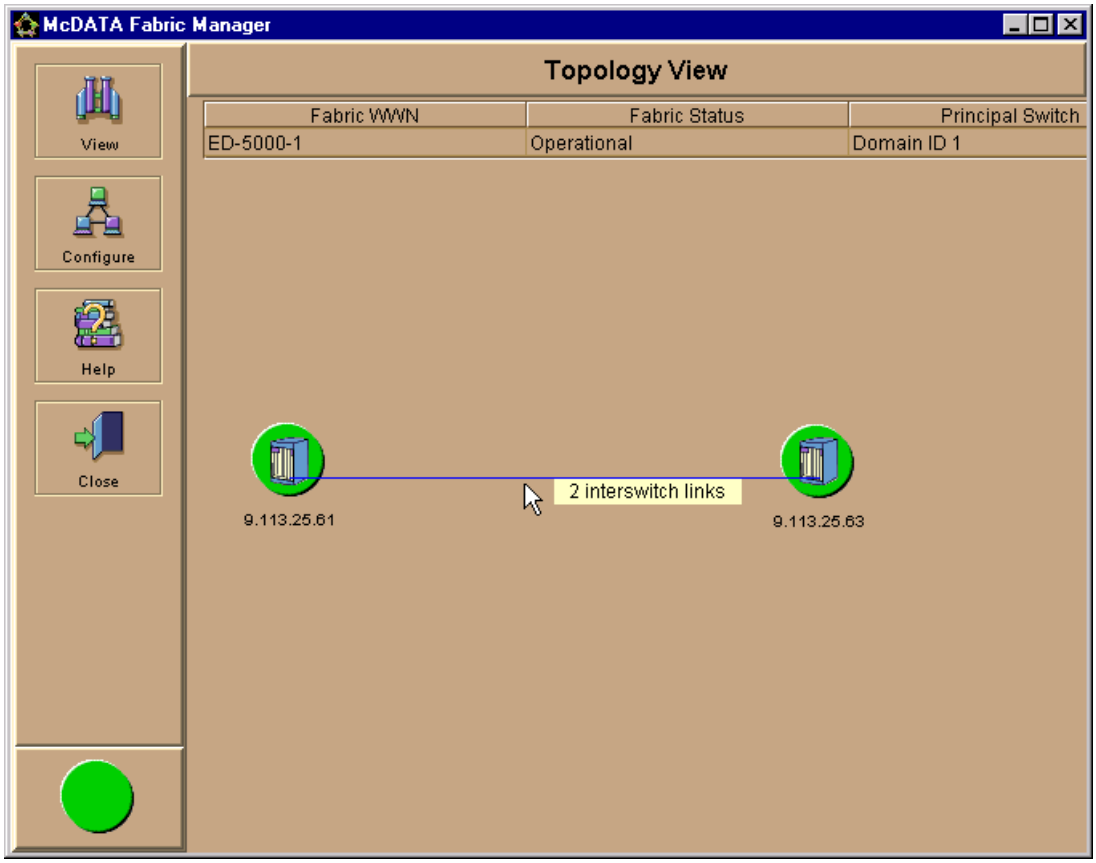

Figure 207. Topology View with two switches connected and configured

This indicates that the connection of the two directors consists of two ISLs. Clicking on either icon will open the Product Manager for the associated director. We can also change the description of the icon by right-clicking on the ISL, as shown in Figure 208. Currently, we see the network address displayed.

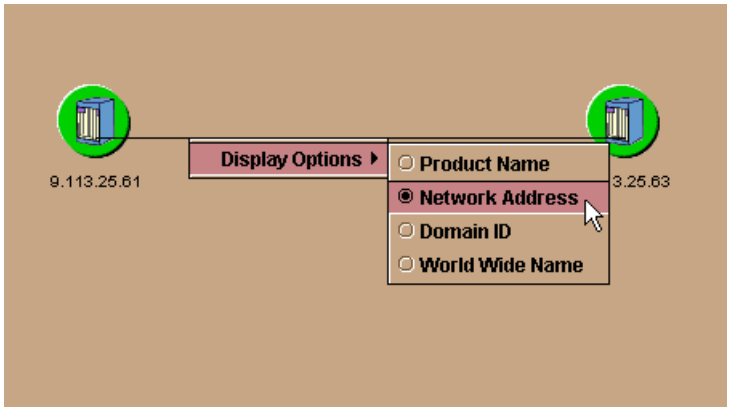

Figure 208. Changing icon text

If the other connected ED-5000 is accessed by means of its own EFC server, the status cannot be determined by the first EFC Server. Therefore, the Topology View looks like Figure 209.

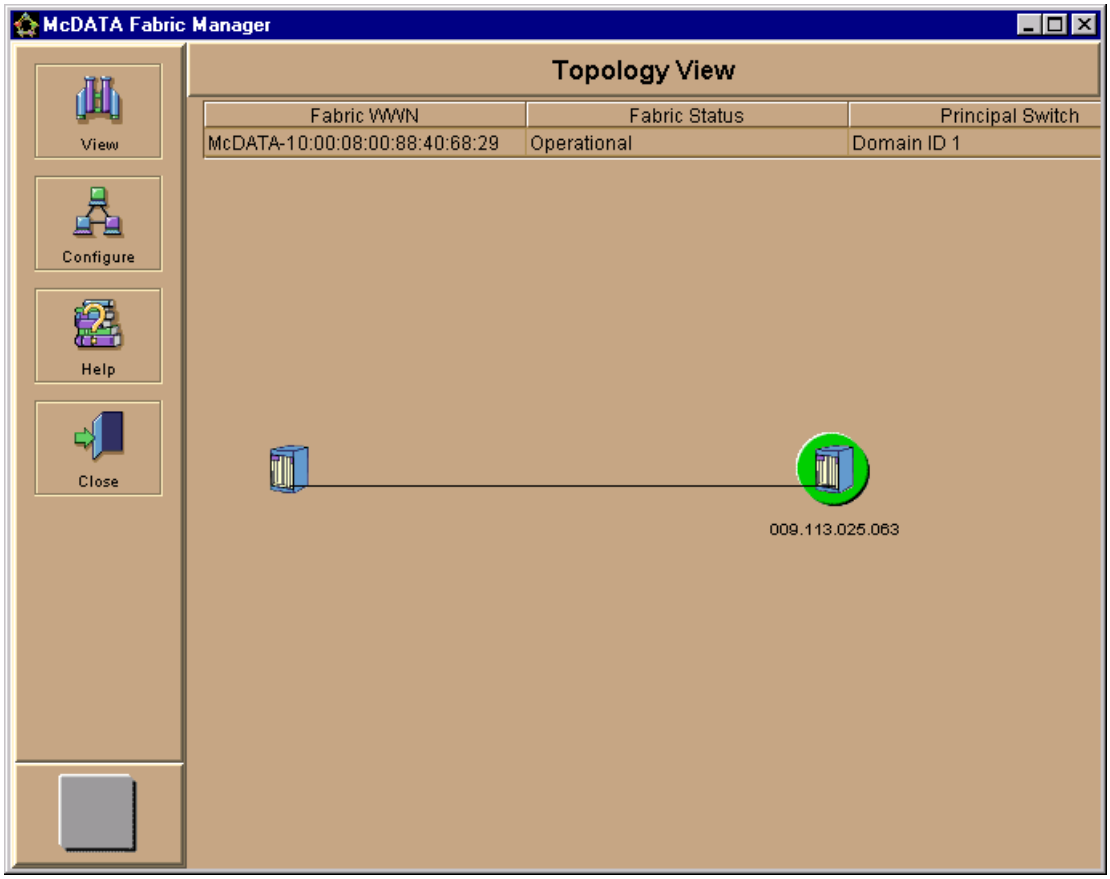

Figure 209. Interconnected ED--5000, one managed by another EFC Server

Changing to the Port List View of the cascaded McDATA, without any other device connected, we see the G_Ports of our two ISLs have turned into E_Ports. We also see the names of the ports we configured, as shown in Figure 210.

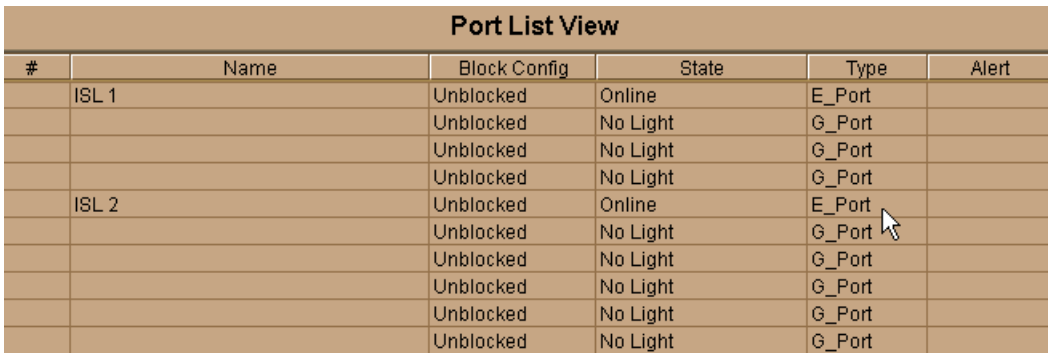

Port List View

#	Name	Block Config	State	Type	Alert
	ISL 1	Unblocked	Online	E_Port	
		Unblocked	No Light	G_Port	
		Unblocked	No Light	G_Port	
		Unblocked	No Light	G_Port	
	ISL 2	Unblocked	Online	E_Port	
		Unblocked	No Light	G_Port	
		Unblocked	No Light	G_Port	
		Unblocked	No Light	G_Port	
		Unblocked	No Light	G_Port	
		Unblocked	No Light	G_Port	

Figure 210. Port List View with two E_Ports

To finish our example, we configure the zoning, and connect our devices to the fabric, as shown in Figure 202 on page 241. After activating the zone set, our Zoning View now looks like that shown in Figure 211.

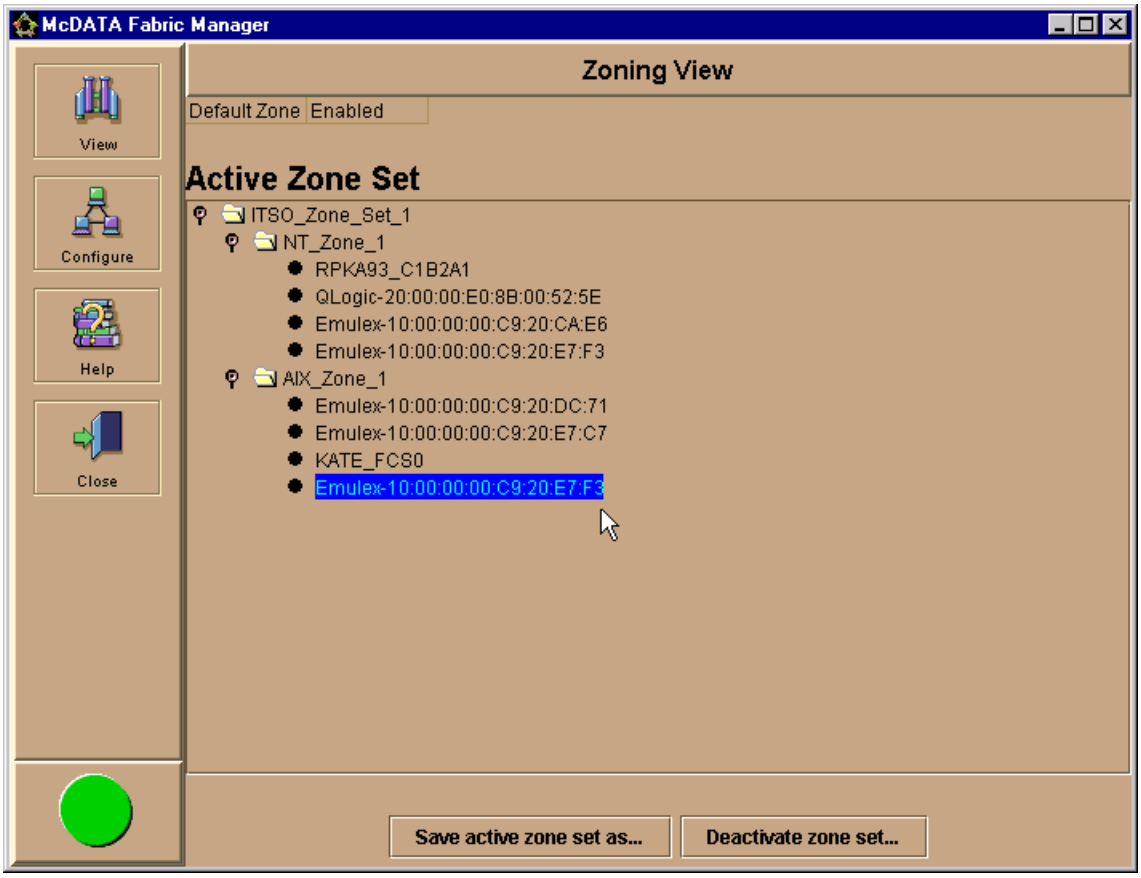

Figure 211. Active Zone Set corresponding to Figure 202 on page 241

We successfully completed all the steps necessary to cascade a McDATA director.

Chapter 8. Implementing the SAN Data Gateway

In this chapter, we describe the steps involved in planning and implementing the IBM Storage Area Network Data Gateway, 2108-G07. The SAN Data Gateway is a hardware solution to allow connection of Fibre Channel ready host systems to attach to SCSI storage systems.

A diagram to show a SAN Data Gateway configuration using a single host is shown in Figure 212.

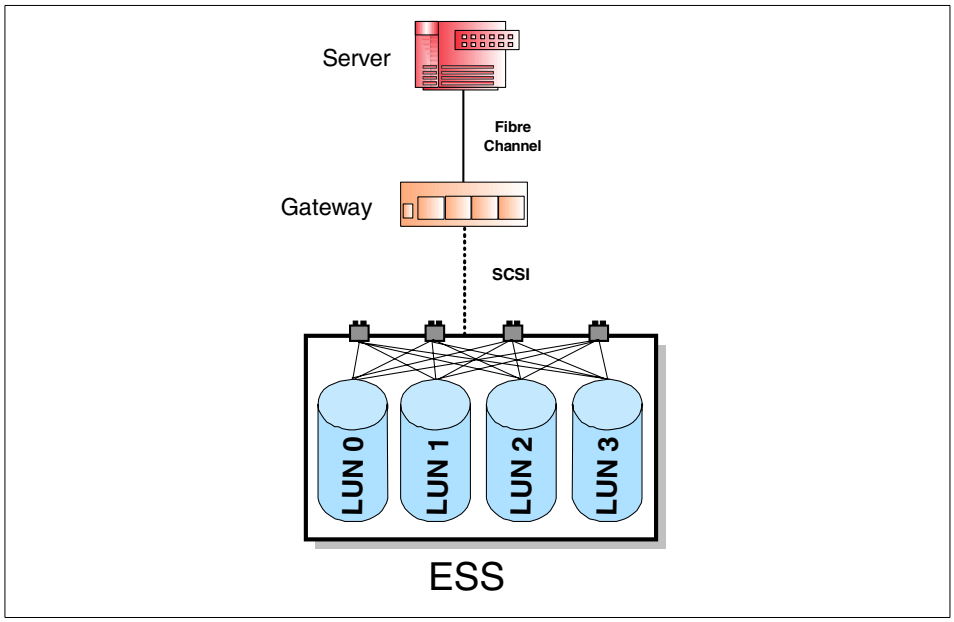

Figure 212. SAN Data Gateway configuration

8.1 SAN Data Gateway

The IBM Storage Area Network Data Gateway is an essential component of the SAN infrastructure. It provides several benefits to bridge the legacy gap as storage products migrate from SCSI based attachments to Fibre Channel.

The IBM Storage Area Network Data Gateway allows you to:

- Protect legacy storage equipment while utilizing the latest host servers with Fibre Channel support

- Expand connectivity to storage devices with use of IBM SAN hubs, switches, and directors
- Perform channel zoning and LUN masking capability to allow access at a volume level
- Overcome the distance limitations of SCSI based host systems using longwave ports that support distances up to 10 km
- Utilize the StorWatch SAN Data Gateway Specialist which is an easy to use interface for managing and controlling access of host systems to storage devices

The SAN Data Gateway is available as a rack-mount unit or as a stand-alone tabletop unit. The gateway model provides two shortwave Fibre Channel ports and four Ultra SCSI Differential ports to attach disk or tape storage devices. One or two Fibre Channel cards – dual-port, shortwave and/or single port, longwave – may be added for a maximum of six shortwave ports, or two shortwave and two longwave ports. If you are using the dual-port shortwave cards, Figure 213 depicts the port assignment numbers for the optical interfaces.

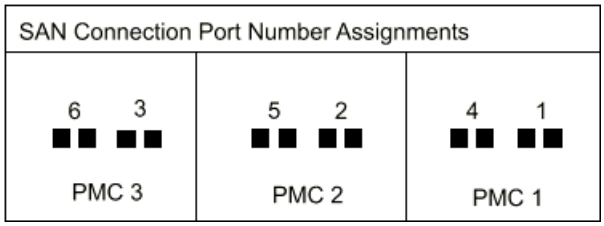

Figure 213. SAN connection port assignment

8.2 Installation

Before any server or device connection is made to the SAN Data Gateway, power on the unit and connect a service terminal to the 9-pin Service port located on the rear right hand side of the unit. A PC running a terminal emulation program, such as Windows HyperTerminal or NETTERM, can be used. The settings of the Service port are:

- 19200 baud
- 8 data bits
- No parity
- 1 stop bit
- X-on/X-off flow control
- VT-100 compatible terminal

If a PC with terminal emulation is used, a 9-pin female to 9-pin female, Null modem cable is required and is provided with the unit.

Once connected, power on the SAN Data Gateway and the start up messages will appear and scroll across the window. When the power on sequence has completed, a prompt `Gateway>` appears on the window as shown in Figure 214.

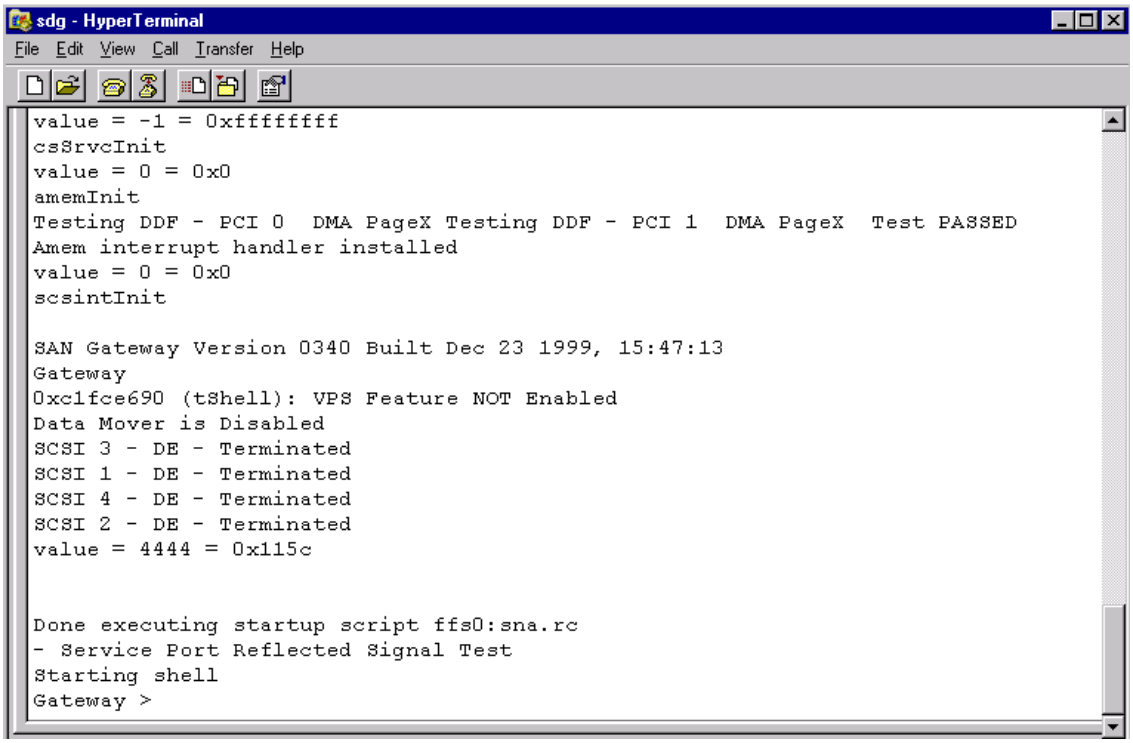

Figure 214. IBM Storage Area Network Data Gateway startup

If you type in `help` and then press Enter, a list of available commands is provided. The commands are case sensitive and *must* be entered as they appear.

Issue the `initializeBox` command to remove any configurations files that may be present. The unit will restart automatically.

Note

The `initializeBox` command will erase all configuration files. It should be used only during initial SAN Data Gateway installation.

8.2.1 Setting the Ethernet address

Once restarted, the Ethernet port must be configured and attached using network information provided by the network administrator. To set the IP address, use the `ethAddrSet` command. The address must contain the double quotes (“):

```
Gateway > ethAddrSet "9.111.24.66"
Network not Enabled
Write complete
Host Address set to 9.111.24.66 for Ethernet interface
value = 0 = 0x0
Gateway > _
```

If a subnet is required then add it after the IP address and separate the two addresses using a comma. For example:

```
Gateway > ethAddrSet "9.111.24.66","255.255.255.0"
```

If a gateway or standard router is to be specified, then issue the `gateAddrGet` command to view if there is one set and the `gateAddrSet` command to set or change it. For example:

```
Gateway > gateAddrGet
No current gateway address set
value = 0 = 0x0
Gateway > gateAddrSet"193.1.1.11"
Write complete
```

The Ethernet port on the SAN Data Gateway comes from the factory disabled. To enable it, you must issue the `ethEnable` command. This will not take effect until the unit is rebooted. The reboot can occur from a power off, or

by issuing the `reboot` command. During the reboot, you will see that the IP address is set and now enabled.

```
Gateway > ethEnable
Write complete
Ethernet will be enabled on next boot
value = 0 = 0x0
Gateway > _
```

8.2.2 Using Telnet on the SAN Data Gateway

If a user would prefer to telnet to the SAN Data Gateway rather than by using the service terminal port after initial setup, this can be done. First you must create a user from the service terminal by using the `userAdd` command. Enter the login name and password using the quotes and comma:

```
ITSO > userAdd "itso","residency"
value = 0 = 0x0
ITSO > _
```

You cannot telnet to the Gateway and use the service port at the same time. When you telnet to the Gateway, the service port on the rear of the unit will stop its communications. After you end the telnet session, then the service port will become available again.

8.2.3 Startup sequence

You must start up the SAN Data Gateway and the attached host and target devices in a specific order. When you add or remove SCSI devices or update firmware, you must restart. The following procedures describe the situations and order of procedure when you restart the SAN Data Gateway.

Before you restart the SAN Data Gateway, you must stop all input and output (I/O) activity between the host and SCSI devices.

1. SCSI devices

   Turn on the SCSI devices. You must turn on all SCSI devices attached to the SAN Data Gateway before you initially turn on or restart the SAN Data Gateway.

2. SAN Data Gateway

   The SAN Data Gateway scans the SCSI buses when it starts. If you add or remove SCSI devices after the Gateway has started, the Gateway will not detect the changes. You can invoke an SCSI rescan or restart operation

from either the StorWatch SAN Data Gateway Specialist client or the service terminal.

3. Fibre Channel host

   Before you turn on or restart the hosts that are connected with Fibre Channel to the SAN Data Gateway, you must wait until the SAN Data Gateway has finished starting. You will know the Gateway has finished starting when the ready light on the front panel blinks at frequency intervals of one second.

   - Some operating systems provide you with software methods that allow you to add or remove SCSI devices dynamically after the host has started. To ensure reliable operation, restart the host.
   - If you update SAN Data Gateway firmware, you must restart the Gateway to use the new firmware. To ensure compatibility between the firmware features or functions and the host, restart the host.
   - If you update SCSI device firmware, the SAN Data Gateway Explorer application does not display the new firmware version until the SAN Data Gateway has issued an SCSI inquiry. The SCSI inquiry occurs when the Gateway rescans the SCSI buses. The SCSI inquiry also occurs when the StorWatch SAN Data Gateway Specialist client application or the service terminal rescans the SCSI buses.

Currently, up to eight different hosts can be attached to each Fibre Channel port. If all six ports are installed, then 48 different hosts can attach to the SAN Data Gateway.

8.3 StorWatch SAN Data Gateway Specialist

The StorWatch SAN Data Gateway Specialistsoftware provides remote capability for all management, configuration, and event notification. It is comprised of three parts:

- Agent
- Server
- Client

Agent

The agent is embedded in the operating system of each SAN Data Gateway to provide a stand-alone manageable host. The StorWatch SAN Data Gateway Specialistsoftware uses SNMP to set and retrieve information that controls the operation of the Agent. The Specialist also uses SCSI over TCP to allow updates to the Gateway and target device.

Server
The server is a Java application that runs on a host and is used to maintain communication with the agents and acts as an intermediary between the agent and the client. The server coordinates the request from multiple clients to manage multiple gateways or agents. Multiple clients can share data the server already knows about, and the server receives all traps from the agent and forwards them to the clients that are registered to receive them.

Client
The client is a Java application that operates from any compatible computer as long as it has a TCP/IP connection to the server. One or more clients can connect to a server. The client provides the user interface to allow the management and configuration of the SAN Data Gateway.

The server and client can be installed on to the same computer.

The StorWatch SAN Data Gateway Specialistsupports the following operating systems:

- Windows 95,98, 2000, and NT 4.0 with SP5 or later
- AIX ver 4.3.3 or later
- Solaris 2.6 or later

8.3.1 Installing StorWatch Specialist

The Specialist software is not bundled with the SAN Data Gateway. The Specialist software is downloaded using a Web browser by going to the IBM Storage Area Network Data Gateway Web site:

`http://www.storage.ibm.com/hardsoft/products/sangateway/support/form1.htm`

This will take you to a registration window. Enter the required information and select `Submit Information`. A license agreement window is shown, and once reviewed, select `I agree`. The Download Main Page window will load. Then select the specific operating system platform. Review the readme.txt file for the latest information and instructions, before installing.

This Web site also contains all the latest firmware for the SAN Data Gateway and supported host bus adapters.

The StorWatch SAN Data Gateway Specialistsoftware file is a self-extracting file. Once it has been downloaded, execute or run the file and it will automatically load onto your computer.

8.3.1.1 Starting the Specialist

To start the Specialist, the server must be started first, and then the client can be launched. Figure 215 provides an example of the StorWatch SAN Data Gateway Specialistwith server and client loaded onto the same Windows NT computer.

Figure 215. StorWatch SAN Data Gateway Specialist startup

Once the Server has been launched you should see a window similar to Figure 216.

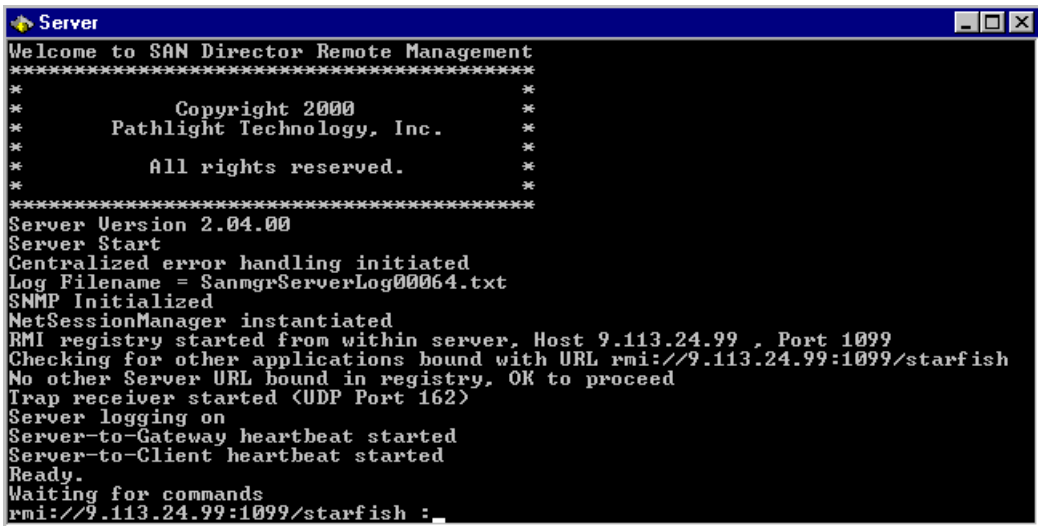

Figure 216. StorWatch SAN Data Gateway Specialist server

The client software can now be launched. If the server and client are not on the same PC, then a dialog box will appear to allow you to enter in the IP address of the computer that has the server software loaded. If the server and client are on the same computer you will be automatically connected to this server. After connection to the server is complete, a dialog box will appear, in which you can enter in a user name and password.

The IBM Storage Area Network Data Gateway provides a default administrator:

- User Name: StorWatch
- Password: StorWatch

The fields are case sensitive so they must be entered in as shown above.

A new administrator account should be set up by selecting **Admin -> Add User** from the toolbar. After a new administrator account is created, then the default user StorWatch is deactivated.

Note

If a new administrator account has been created and the password is lost and no other account has administrator access, a service representative must be contacted.

8.3.2 Using the StorWatch SAN Data Gateway Specialist

Once you are logged in to the Specialist, you must now connect to the SAN Data Gateway. A dialog box appears requesting the IP address of the SAN Data Gateway. As it connects, it will download the information from the SAN Data Gateway and be presented on your window.

If a dialog box does not appear automatically, select **Tools-> Connect SAN Data Gateway or SAN Data Gateway Router** from the toolbar. This can also be used to connect to several Gateways or Routers from a single client session.

In Figure 217, we show the initial view once a connection to a SAN Data Gateway is established and the data has been downloaded.

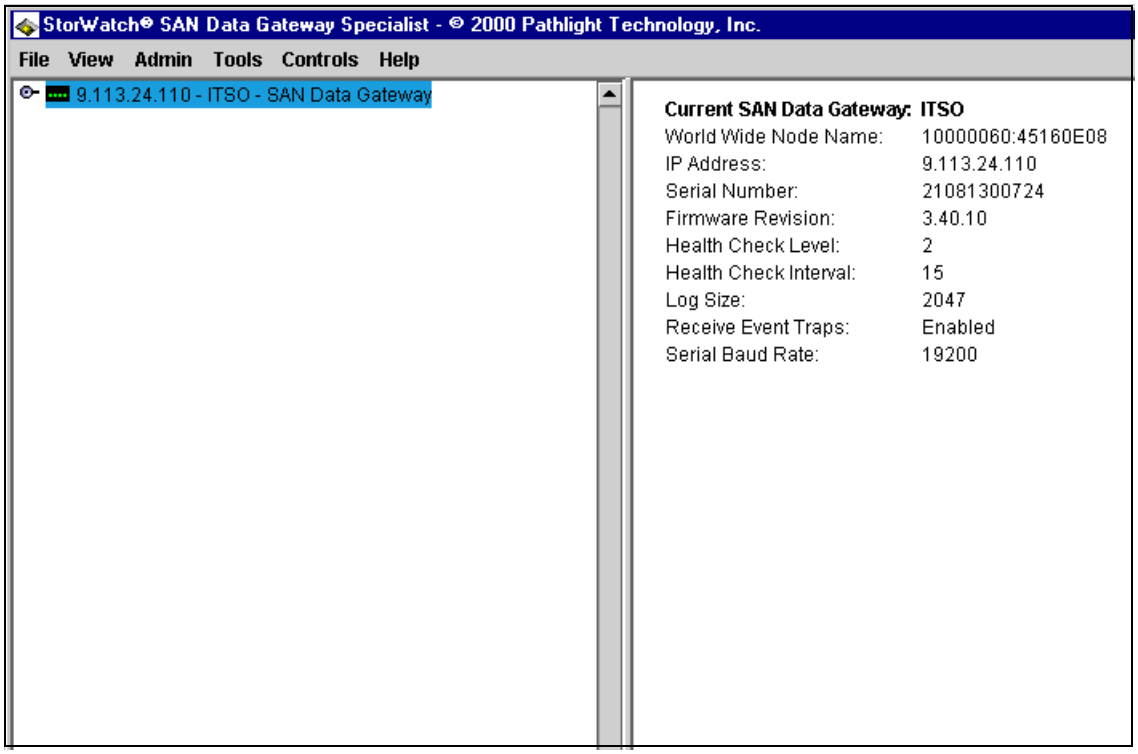

Figure 217. StorWatch SAN Data Gateway Specialist initial view

The left side of the window shows the SAN Data Gateway unit we are connected to and the right side provides product data information. You will also notice that the toolbar will have options available that were previously greyed out. You can now connect to another SAN Data Gateway, disconnect from a SAN Data Gateway, enable and access the Zoning and VPS features, restart the Gateway, and also refresh the data to your window by downloading it again.

These options become available when a SAN Data Gateway is highlighted. As you begin to add SAN Data Gateway systems or drill-down into a particular Gateway by selecting and highlighting different channels or ports, different options will become available and other options will become greyed out and unavailable. Be aware of what system, channel, or port is highlighted as you move through the toolbar.

As we can connect to several SAN Data Gateway systems from one client session, select the particular Gateway you want and it will be highlighted in blue as shown in Figure 218.

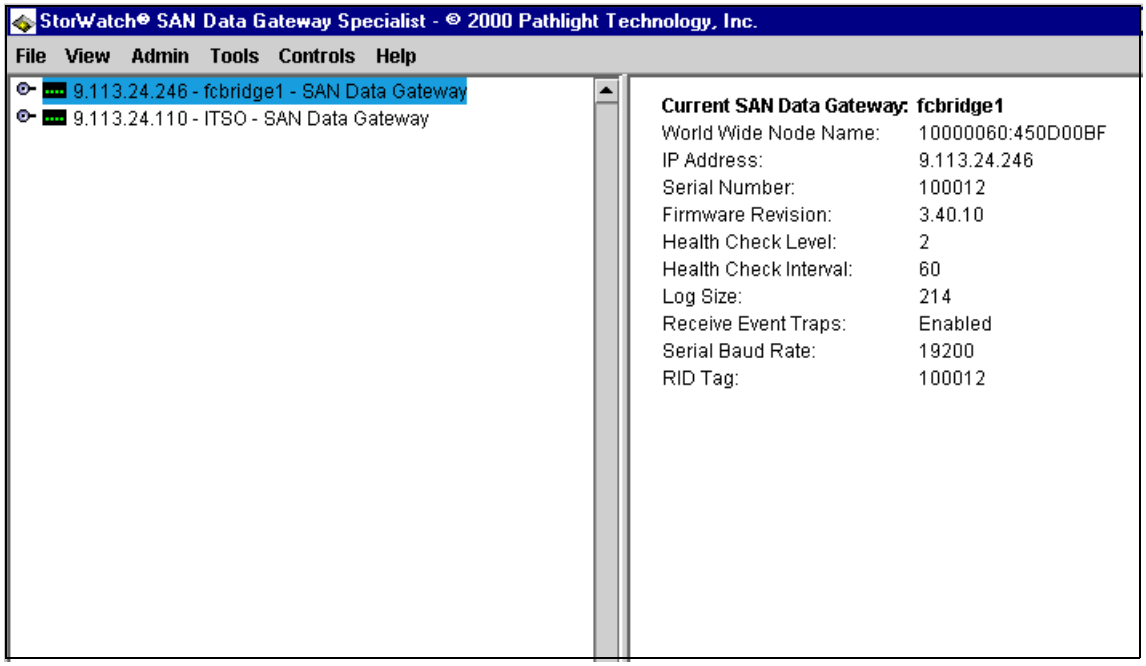

Figure 218. Selecting from multiple SAN Data Gateways

On the left hand side of the highlighted Gateway, there is a small key, and by selecting this, it expands the view to show you all SCSI ports and installed Fibre Channel ports. Figure 219 shows a Gateway with four SCSI ports and two Fibre Channel ports.

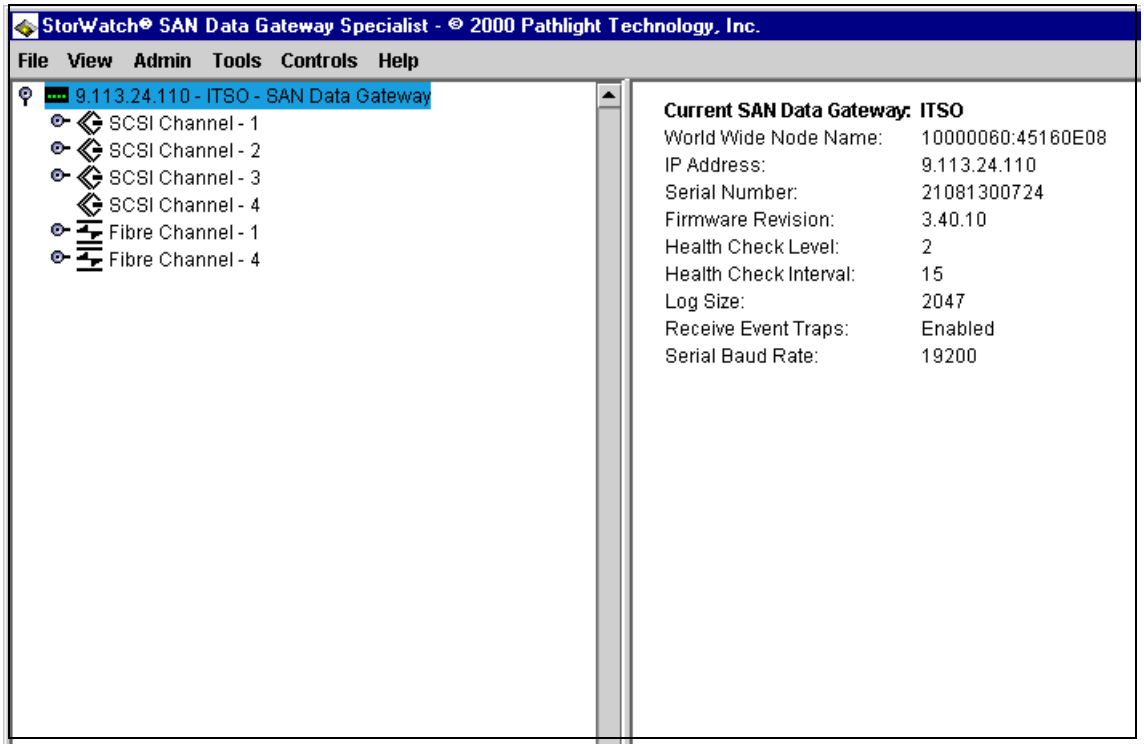

Figure 219. Expanded Gateway view

SCSI channels 1, 2, and 3 and Fibre Channel ports 1 and 4 also have a key on the left hand side to depict that there are devices attached. By selecting a key, you will now expand the tree, as seen in Figure 220, and view the different disk devices attached.

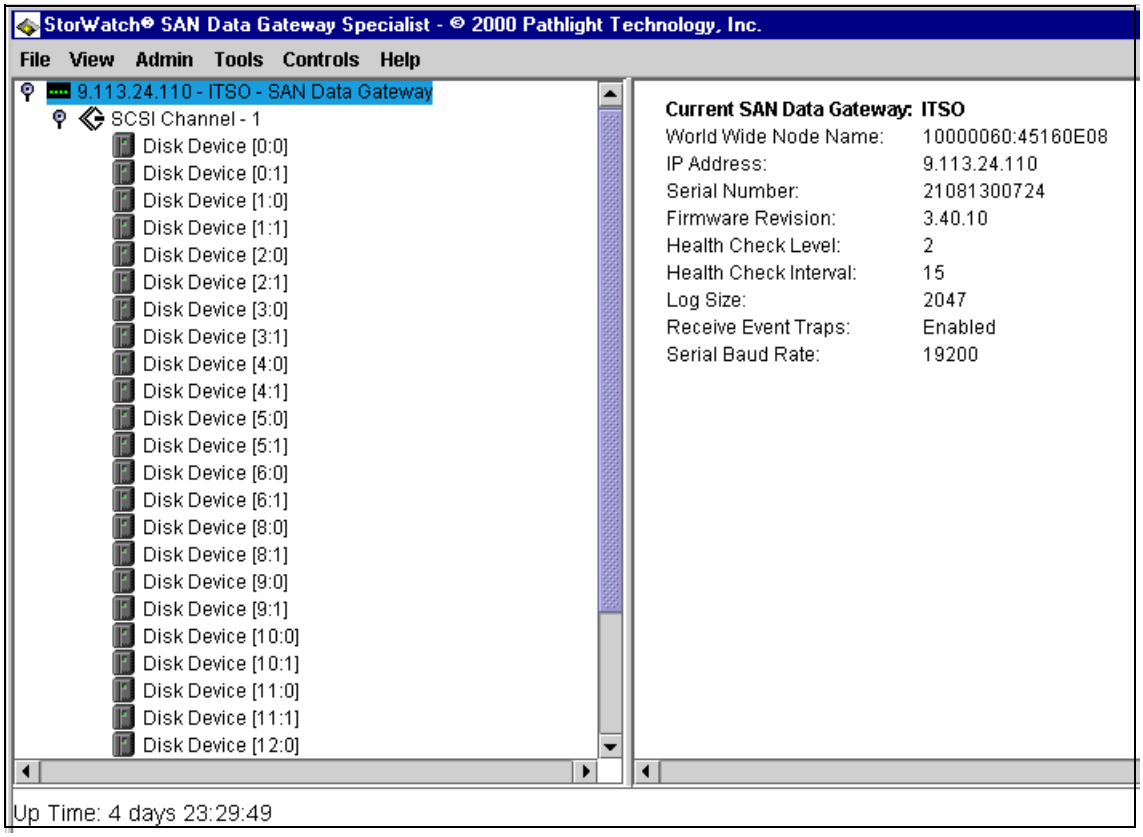

Figure 220. SCSI channel expanded view

You can also select and highlight each SCSI Channel. You will notice that as you do this, the information window on the right side will provide data that is unique to that SCSI channel, as shown in Figure 221.

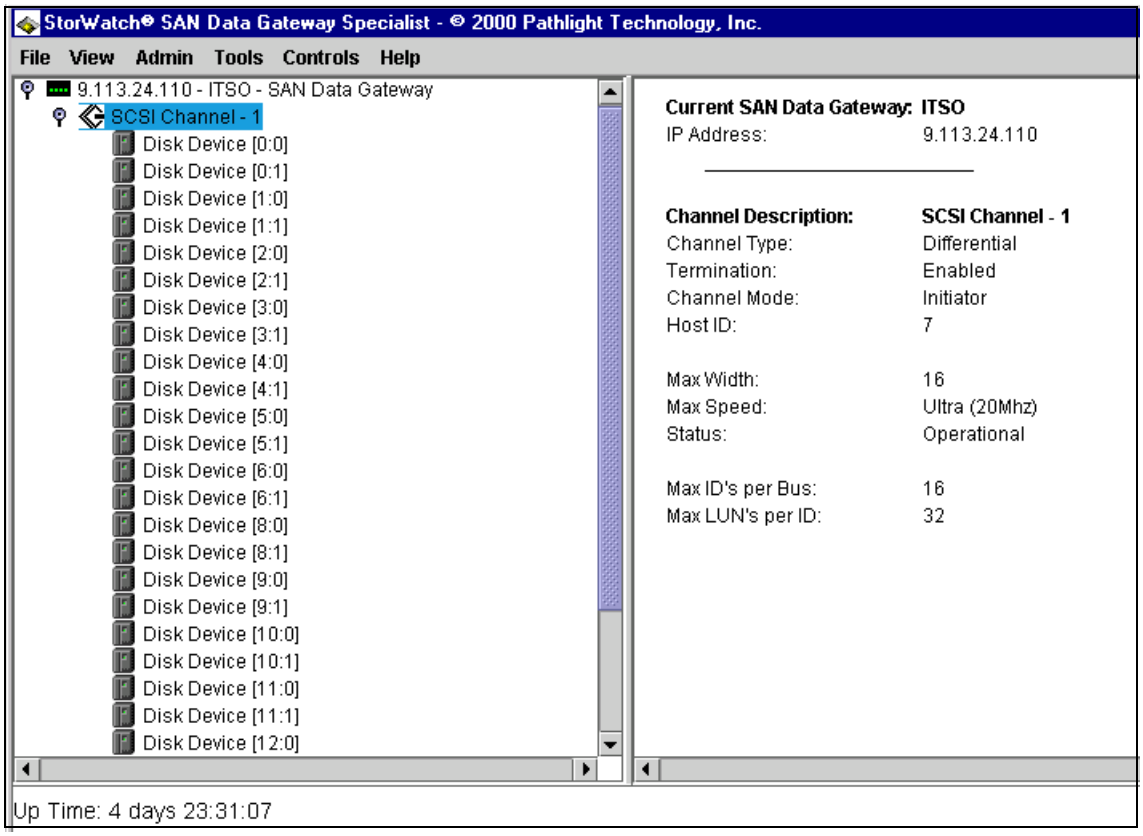

Figure 221. SCSI channel data

Information pertaining to a particular disk device is shown in Figure 222.

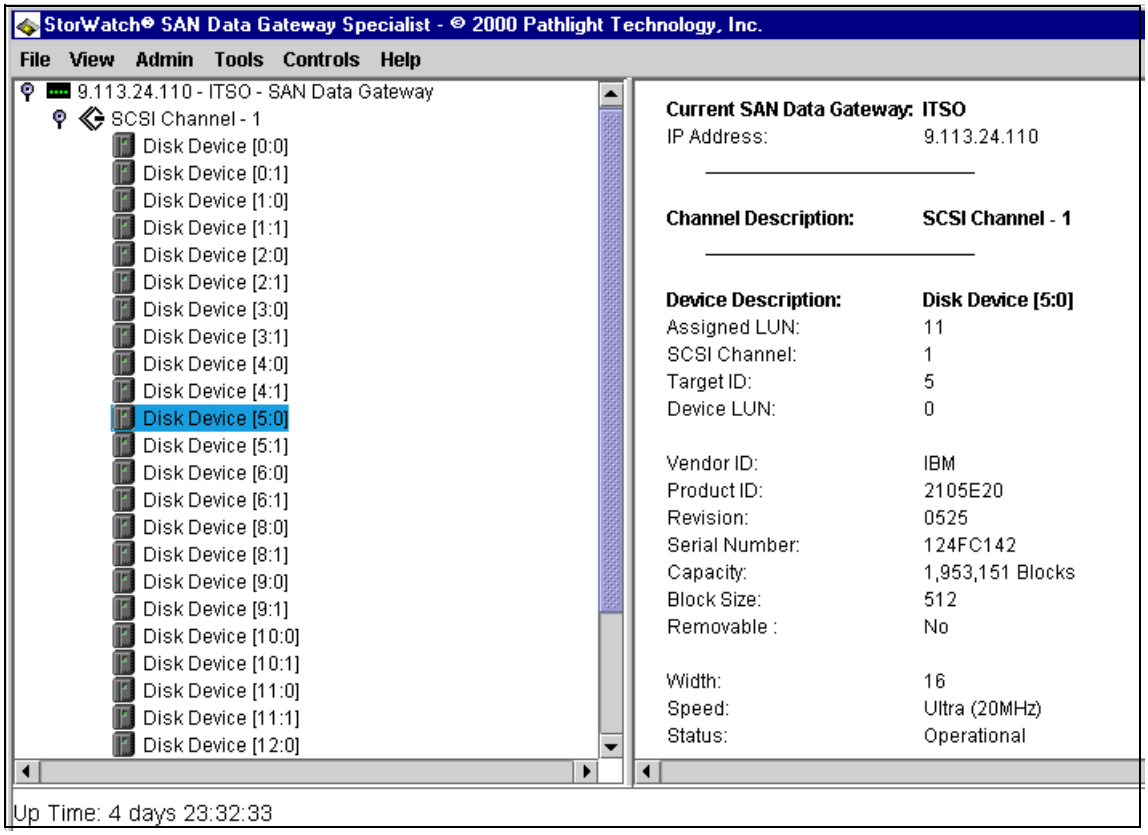

Figure 222. Disk device data

You can perform the same drill-down on the Fibre Channel host as we did with the SCSI channel and disk devices. Select one of the Fibre Channel port connections, as shown in Figure 223, and its data will be shown on the right hand view pane.

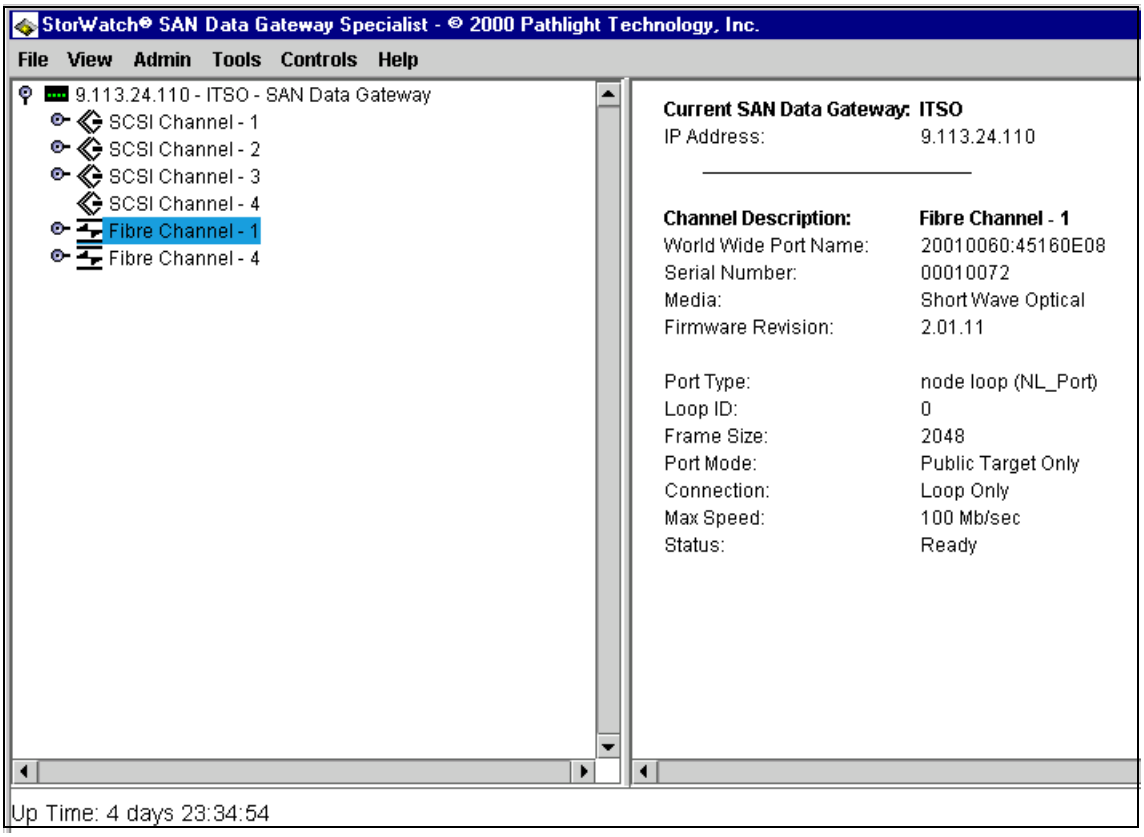

Figure 223. Fibre Channel port data

By selecting the key to the left, you can expand the tree and select the host system attached to that port. Figure 224 shows the detail on the specific host.

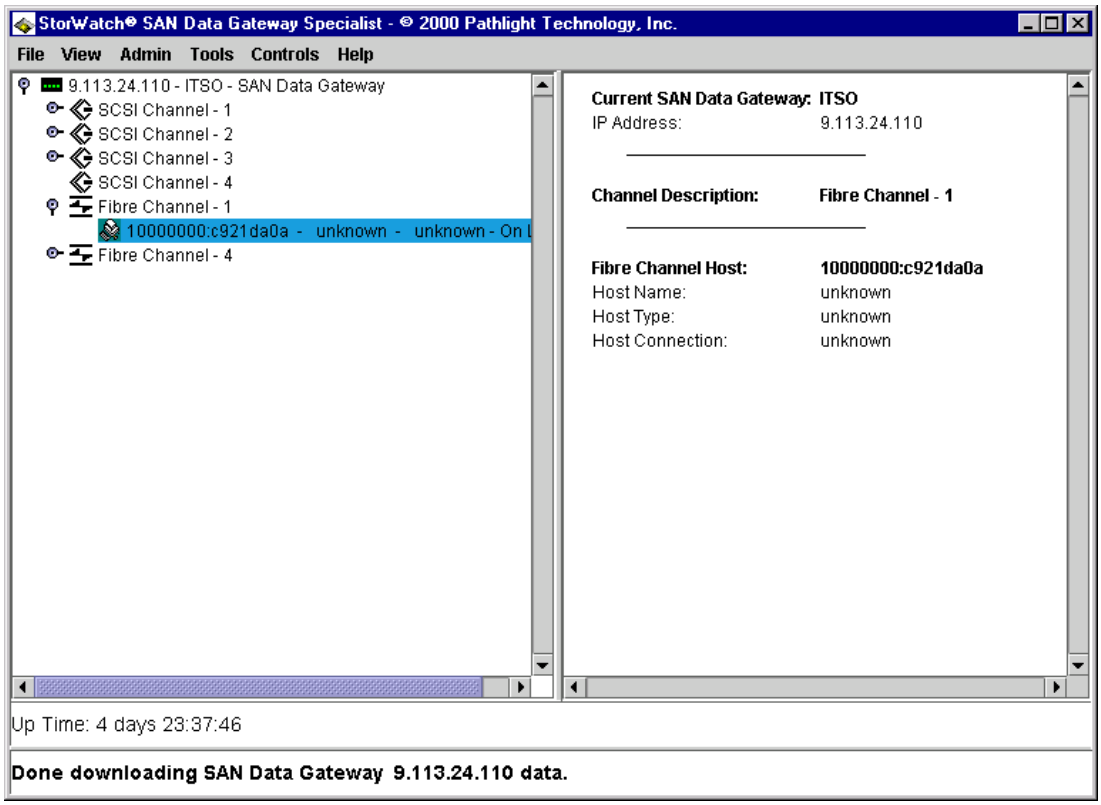

Figure 224. Fibre Channel host data

As you select and highlight the different ports or devices, there are different options available from the top toolbar. If an SCSI channel is highlighted, select **Controls** from the toolbar. You will notice that all options are grayed out except for **SCSI Channel**. Once selected, a dialog box will appear, as shown in Figure 225, and display the settings for the SCSI channel.

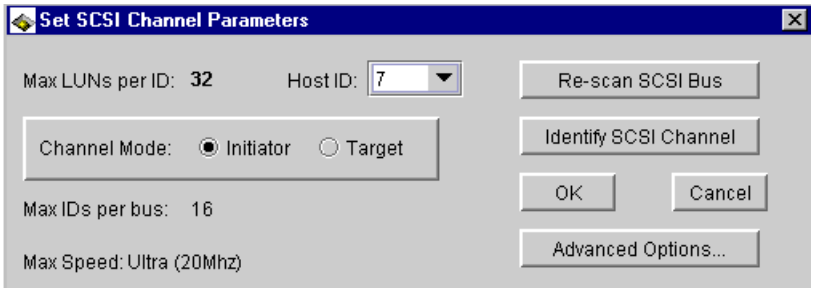

Figure 225. SCSI channel parameters

Selecting the **Advanced Options** button displays a dialog box, as shown in Figure 226. These settings are not typically changed and may disrupt normal operations. Refer to a service representative before changing any Advanced Options.

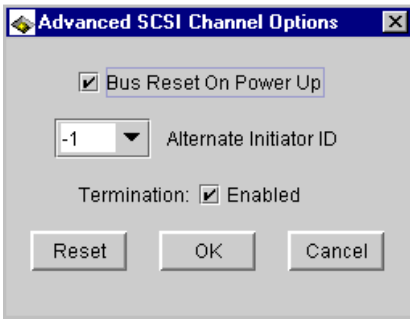

Figure 226. Advanced SCSI parameters

Fibre Channel parameters are displayed in a similar fashion. Highlight a Fibre Channel port and select **Controls** from the toolbar, you will notice that now all options are grayed out except the **Fibre Channel** option. By selecting this a dialog box will display the parameters that can be changed for the Fibre Channel port selected. If any of the settings, as shown in Figure 227, are changed, the SAN Data Gateway must be restarted.

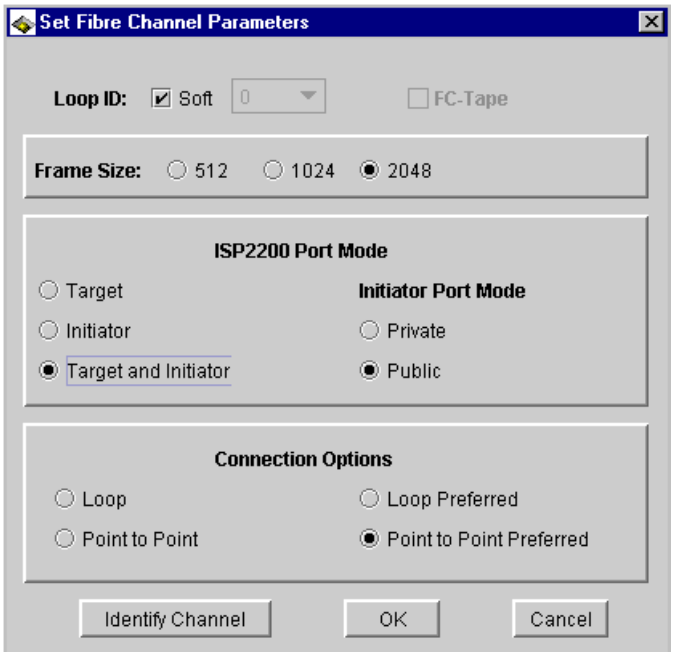

Figure 227. Fibre Channel port parameters

A detailed description of the SCSI and Fibre Channel settings can be found in the *IBM Storage Area Network Data Gateway Installation and User's Guide*, SC26-7304.

8.4 SCSI Devices

The four SCSI ports on the SAN Data Gateway support Differential Ultra Wide SCSI devices. It will automatically negotiate speed for wide or narrow bus width devices as well as standard, fast and, ultra speeds. The SAN Data Gateway provides a termination on each of its SCSI buses. The termination can be separately disabled if so desired from the StorWatch SAN Data Gateway Specialist window.

8.4.1 LUN support

The SAN Data Gateway provides support for up to 256 LUNs. Each SCSI channel supports up to 15 SCSI targets and up to 32 LUN per target. This is subject to the overall total of 256 available. The first LUN (LUN 0) is used for the Gateway for control and command purposes. That leaves 255 allowable LUN addresses to be used for attaching SCSI devices to the four SCSI ports.

Note

Do not attach more than 255 LUNs to the SAN Data Gateway. Exceeding the 255 limit will cause unpredictable results.

The SCSI devices must be previously set up in your host prior to attaching to the SAN Data Gateway. For example, if an IBM Enterprise Storage Server is to be used as the target device, then all the volume or LUN assignments must be completed. Also, the volumes should be assigned to the appropriate SCSI port in the ESS. Attach the SCSI cable from the SCSI device to a SCSI port on the Gateway.

8.4.2 Device discovery

Once attached, the SAN Data Gateway can be restarted or powered on to perform a discovery of the attached devices. If this was done after the `initalizeBox` command from the service terminal, then the initial discovery will be carried out in a specific order. The buses are scanned in order from 1 to 4 and each bus is scanned from SCSI target 0 to 15 and LUN 0 to 32 for each ID. As a device is discovered it is assigned a specific LUN number by the SAN Data Gateway. This numbering will begin at LUN number 1, zero is reserved for the SAN Data Gateway control, and continues sequentially as devices are discovered.

The SAN Data Gateway must add this extra layer of addressing as the host is no longer directly attached to the disk devices but will see a single target ID, the SAN Data Gateway. In a regular SCSI environment with the host attached directly to the SCSI device, the host is able to recognize the target and LUN ID of that device. Since we add the SAN Data Gateway in between the host and the device, the host is no longer directly attached to the SCSI device to recognize its target or LUN ID.

8.4.3 Assigning LUN IDs

The Gateway assigns LUN IDs to all the attached devices, up to the maximum of 255, and it creates a map of the actual disk device target and LUN ID to its own assigned LUN ID. This map is stored in nonvolatile memory within the Gateway.

8.4.4 Adding SCSI devices

New SCSI devices can be added at any time. If a new device is added, for example, another volume in the IBM Enterprise Storage Server has been

assigned to a SCSI port on a San Data Gateway, the SCSI buses must be rescanned to detect the new device. This can be done from the service terminal using the `scsiRescan` command or from the StorWatch SAN Data Gateway Specialist. If using the Specialist, select and highlight the SCSI bus that the device has been assigned to, and select **Controls-> SCSI Channel-> Re-Scan SCSI Bus**. As it rescans, the SAN Data Gateway will assign the next available LUN address to the new device. Refresh the data on your Specialist window by selecting **View-> Refresh SAN Data Gateway**.

Once a device has been assigned a LUN ID by the Gateway, it will maintain that ID since it was written into the device map. This is useful in case the device is lost or needs to be replaced. Remove the old device, set the new device to the same SCSI bus target and LUN as the old device and attach it to the same channel. You must rescan the SAN Data Gateway SCSI bus for it to update its data. You will notice that the replaced device has kept the same assigned LUN.

If a device is no longer attached or no longer required, then the assigned LUN is also no longer required. To free up this assigned LUN ID you can issue the `mapWinnowDatabase` command from the service terminal. The existing devices will maintain their assigned LUN IDs.

If a complete rebuild of the SCSI device map is desired or required, this can be done from the service terminal only. You would issue the `mapRebuildDatabase` command. This command deletes the existing device map and the SAN Data Gateway will restart. When it is rebooted, a new scan of the SCSI buses is done as if the system was brand new and assign LUN IDs as described earlier.

When this command is issued, you will also be given an option to clear the Virtual Private SAN (VPS) access settings. Because this allows host access to specific LUNs and by issuing this command, we may change the assigned LUN ID; it is recommended that you always say `yes` to this option. If not, a host may access a volume that you do not want it to access and be restricted from a volume that it had access to previously. The system administrator must rebuild the VPS settings, if enabled, to allow host access to the desired LUNs.

8.5 IBM Storage Area Network Data Gateway access options

The SAN Data Gateway includes two features that are helpful in providing control and security of host access to SCSI devices:

- Zoning
- Virtual Private SAN (or LUN masking)

8.5.1 Zoning

Channel zoning is a feature included with the IBM Storage Area Network Data Gateway. Zoning allows you to restrict access between SAN Fibre Channel connections and SCSI channels. The default settings allow all SAN connections to access all SCSI channels.

8.5.1.1 Configuring zoning

To configure zoning, and you must be an administrator, select **Control -> Access Options -> Channel Zoning** from the toolbar. Figure 228 shows a zoned SAN Data Gateway. A check mark will allow access, and in Figure 228 you can see that SAN connection 1 has access to SCSI channels 3 and 4, but not to SCSI channels 1 and 2. SAN connection 4 has access to SCSI channels 1 and 2, but not to SCSI channels 3 and 4. To change the settings, click on any box and the check mark will toggle on and off. All combinations are possible. Once the desired settings are selected, click **OK**. For the new zone settings to take effect, the SAN Data Gateway must be restarted.

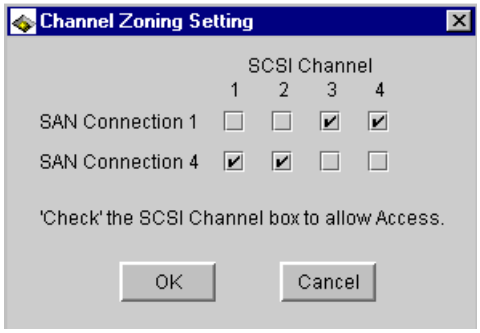

Figure 228. IBM Storage Area Network Data Gateway channel zoning

8.5.2 Virtual Private SAN

The IBM Storage Area Network Data Gateway also provides LUN masking through a feature called Virtual Private SAN (VPS). This provides a granular restriction of host access to specific LUNs while utilizing and sharing the same connectivity paths through the SAN Data Gateway.

VPS keeps track of Fibre Channel hosts by using their unique World Wide Name. In this way, if a switch is attached, the Gateway will also keep track of it, because the switch also has a unique World Wide Name.

VPS also includes a host registration service so that when a host is attached, its unique WWN is registered to the Gateway to provide information on the host and its status.

8.5.2.1 Enabling VPS

The Virtual Private SAN feature has to be enabled, because it comes disabled from the factory. To enable it, make sure the desired SAN Data Gateway is highlighted, if more than one appears in your Specialist window. Select **Controls -> Feature Enable -> Virtual Private SAN**, as shown in Figure 229. A dialog box appears requesting the entry of a license key number. For units with serial numbers lower than 1300600 you are required to get a license key that is matched to the serial number of the SAN Data Gateway. For units with serial numbers higher than 1300600, the VPS feature is bundled into the product. Type the word `enable` and the feature will now be enabled.

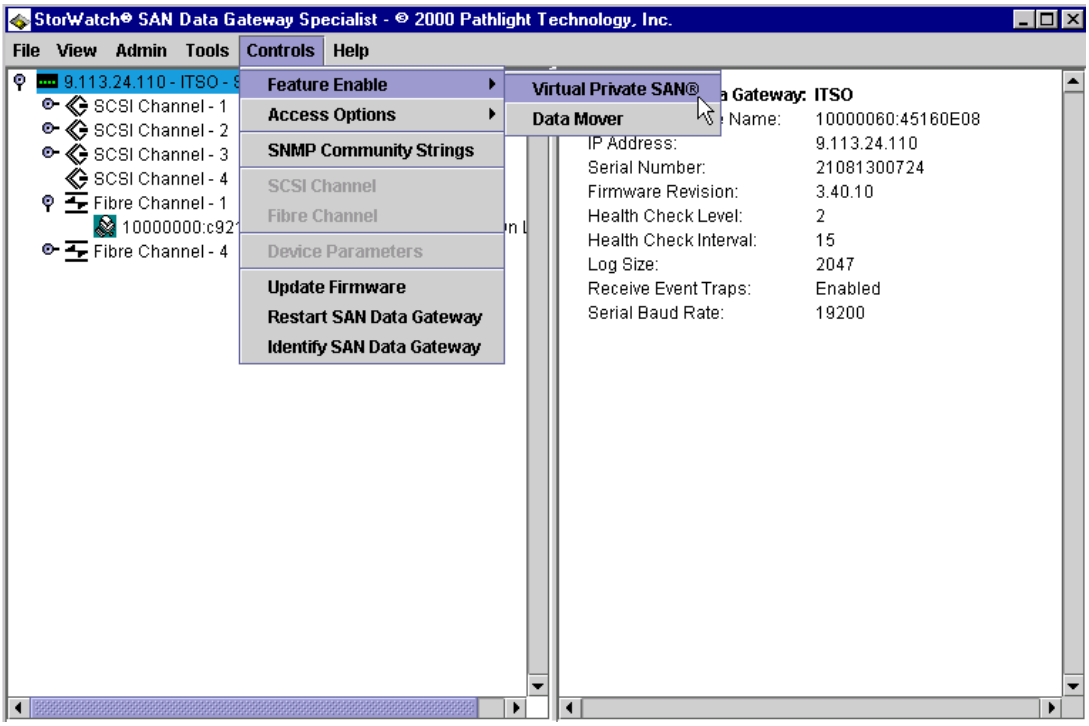

Figure 229. Enabling Virtual Private SAN

8.5.2.2 Host registration

With VPS enabled, the host registration service is also available. This service is provided to simplify VPS configurations. It provides the host name, host type, host connection, and the unique WWN of the Fibre Channel host bus adapter. Host information is sent over the Fibre Channel connection to the Gateway.

To obtain this information, a program is loaded on each host. This program is found, and can be downloaded at the same Web site that the StorWatch SAN Data Gateway Specialist was downloaded from.

`http://www.storage.ibm.com/hardsoft/products/sangateway/support/form1.htm`

After completing the registration and license agreement, the Download Main page is displayed. Select the operating system software subheading and look for “HOSTSW”. View the read.me file for the latest information. Download the software and install it onto the host that will be attaching to the SAN Data Gateway. Follow the instructions provided in the read.me file.

For Windows NT, the file is a self-executing file, so it can be executed or run as any *.exe file. As it runs through the install, make sure to select SAN Explorer and VPS Registration Service, as shown in Figure 230.

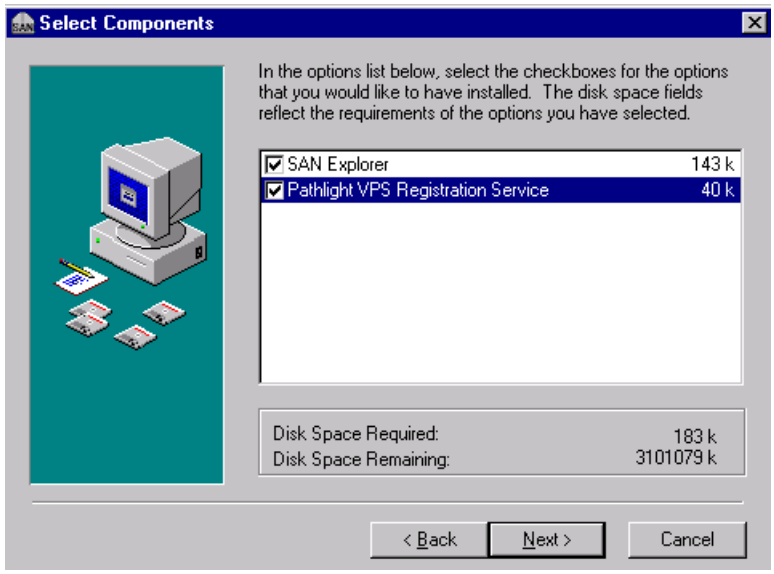

Figure 230. Loading VPS Registration software on Windows NT

Once installed the service runs automatically and does not require further administrator interaction.

8.5.2.3 Communicating with the Gateway

Once the host is attached to the Gateway and restarted, the registration service will communicate to the Gateway. The data shown on the Gateway will have to be refreshed by selecting **View-> Refresh SAN Data Gateway**. This will cause the updated data to be shown on the Specialist window.

The registration service will re-register the host to the SAN Data Gateway at a default of 15 minute intervals. This interval can be changed if so desired.

Previous to enabling the VPS feature, you will have seen the Specialist display a window similar to Figure 231. There is no key beside the Fibre Channel ports to indicate that no host is recognized.

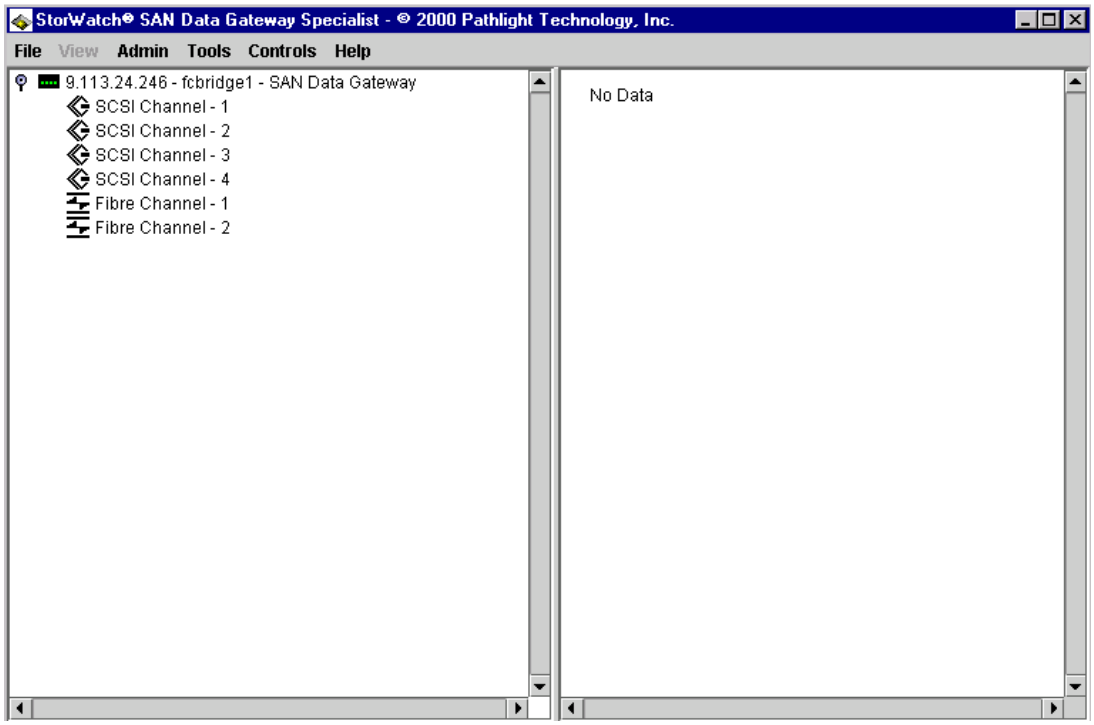

Figure 231. Specialist display without VPS enabled

After the VPS is enabled and a host has registered with the Gateway, all its information will load automatically into the database of the Gateway. Figure 232 shows, in the right hand side view pane, the host name, host type, HBA and connection information, and the unique WWN that was obtained automatically by host registration.

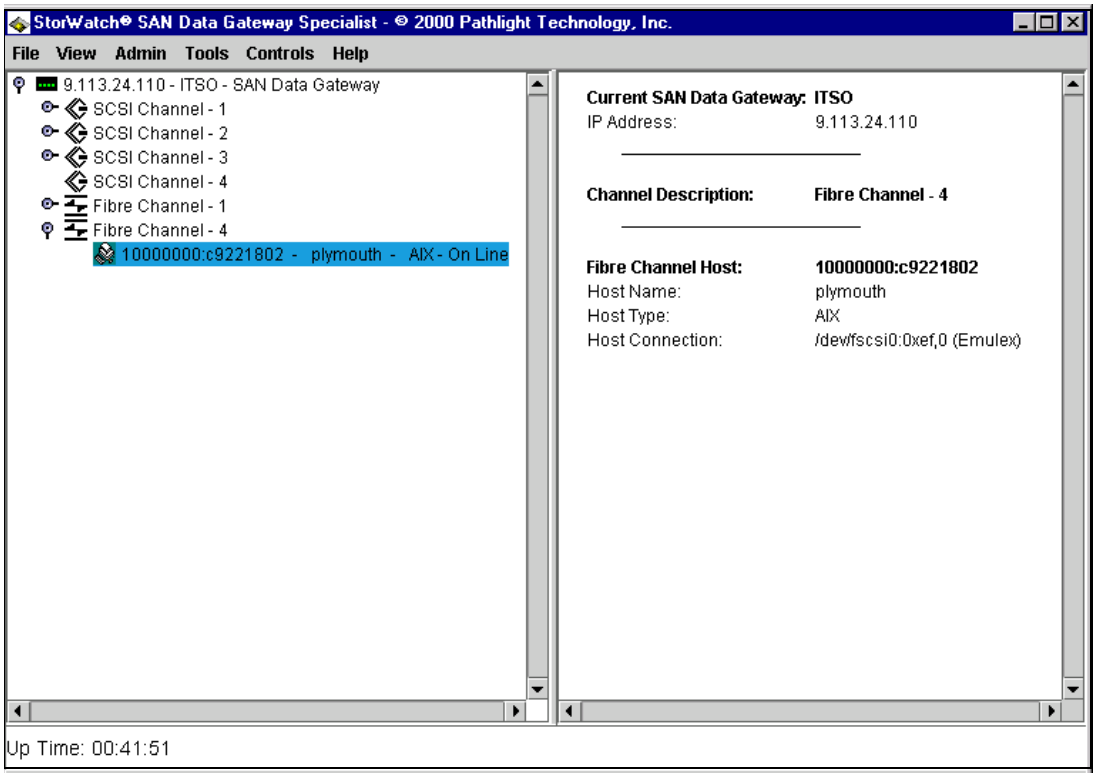

Figure 232. Specialist after VPS enabled and host registration

If no registration software is loaded on the host, or is not available for a specific operating system, for example a Fibre Channel switch, only the WWN of the attached system will register to the VPS database, all other fields will have *unknown*. This is shown in Figure 233.

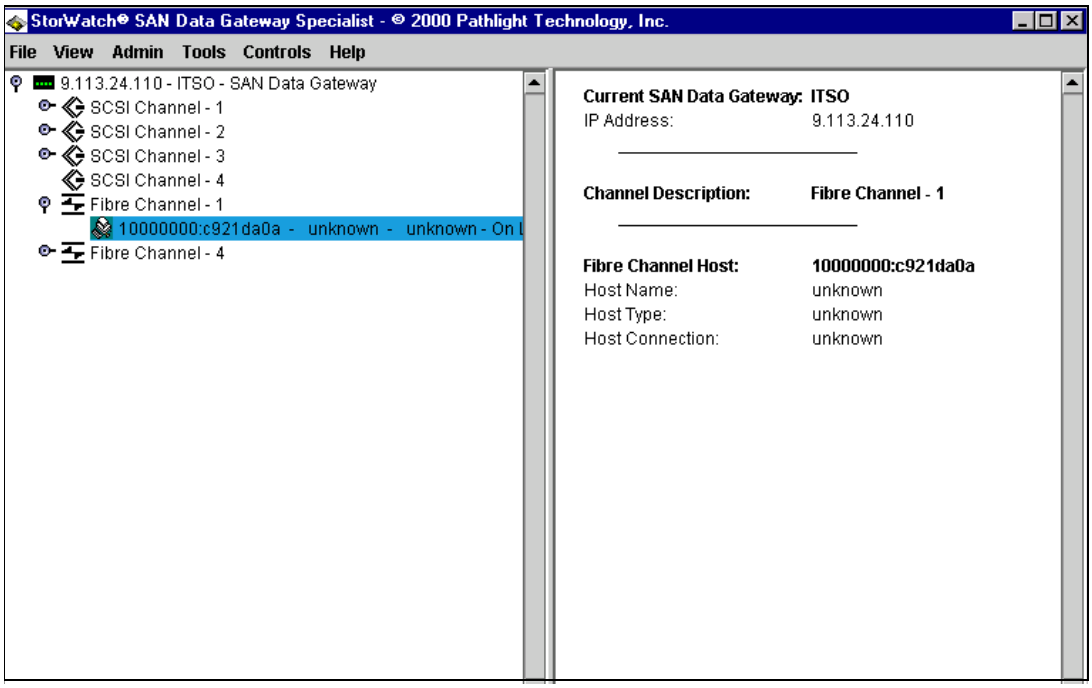

Figure 233. Host system with no host registration software

A diagram to depict the setup that was described, is shown in Figure 234. Note that as the first dual Fibre Channel port on the Gateway is used, it is numbered as input 1 and 4. A description on the Fibre Channel port numbering is shown in Figure 213 on page 252.

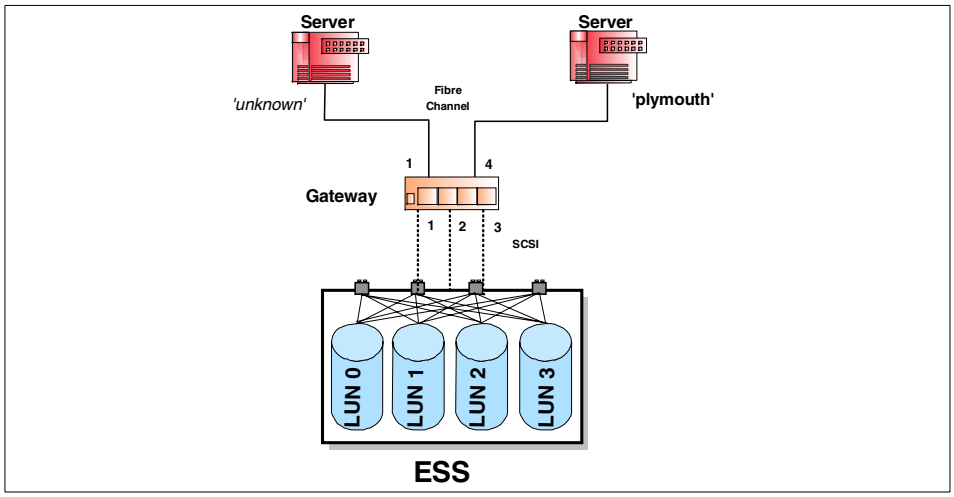

Figure 234. SAN Data Gateway with two hosts

8.5.2.4 Adding host and connection information

The host name and host type and connection information can be added manually by selecting the VPS feature and modifying the data. To do this, select the desired SAN Data Gateway so that it is highlighted. Select **Controls -> Access Options -> Virtual Private SAN**, and you will enter into the VPS settings window. The hosts and their information will appear on the left hand side. Select the host to be modified, and the information is transferred to the bottom part of the window where the modifications can take place, as shown in Figure 235.

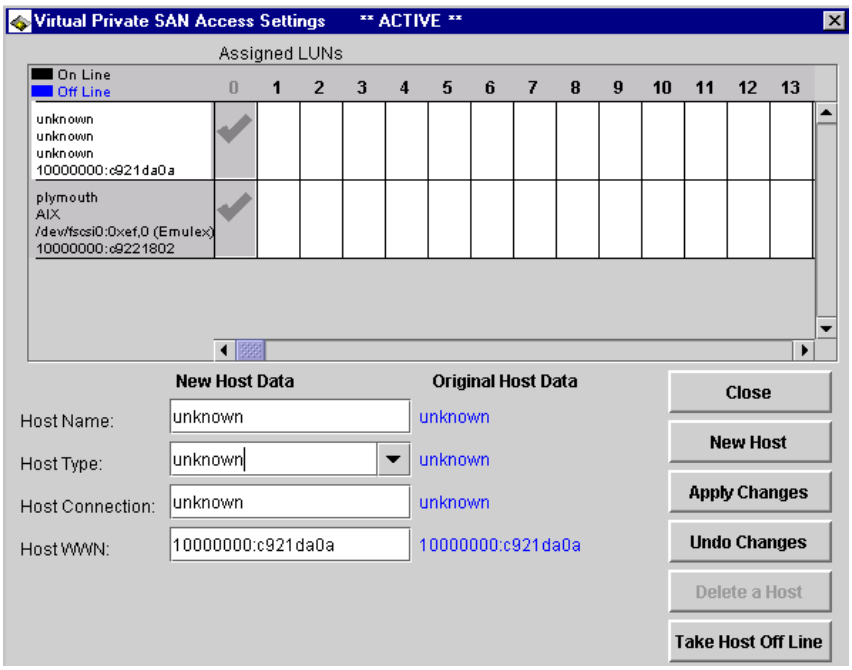

Figure 235. VPS host settings

The Host information can now be entered in the left hand column labeled New Host Data. The Original Host Data column displays the values before any changes are made. Once completed, select **Apply Changes** and then **Close**.

This window can also be used to enter in host information for a host that has not yet connected to the SAN Data Gateway. Select **New Host** and then enter in the WWN and other parameters that are available. You *must,* at minimum, enter in a WWN, and when entering in the number, a colon ':' *mus*t be used to separate the first four bytes from the last four bytes.

Also, you will notice that the host description in the top can change color. If the information is in blue, then the host is offline. If the host information is in black, then it is online.

The VPS feature allows an administrator to quickly view the host information and status at a central location.

8.5.2.5 Setting up a Virtual Private SAN

Remember that the SAN Data Gateway assigns its own LUN numbers and in a sequential order. An attached disk device may have a SCSI target of 1 with a LUN ID of 4, but when recognized by the SAN Data Gateway, its LUN

number assigned will be something completely different. An example is shown in Figure 236.

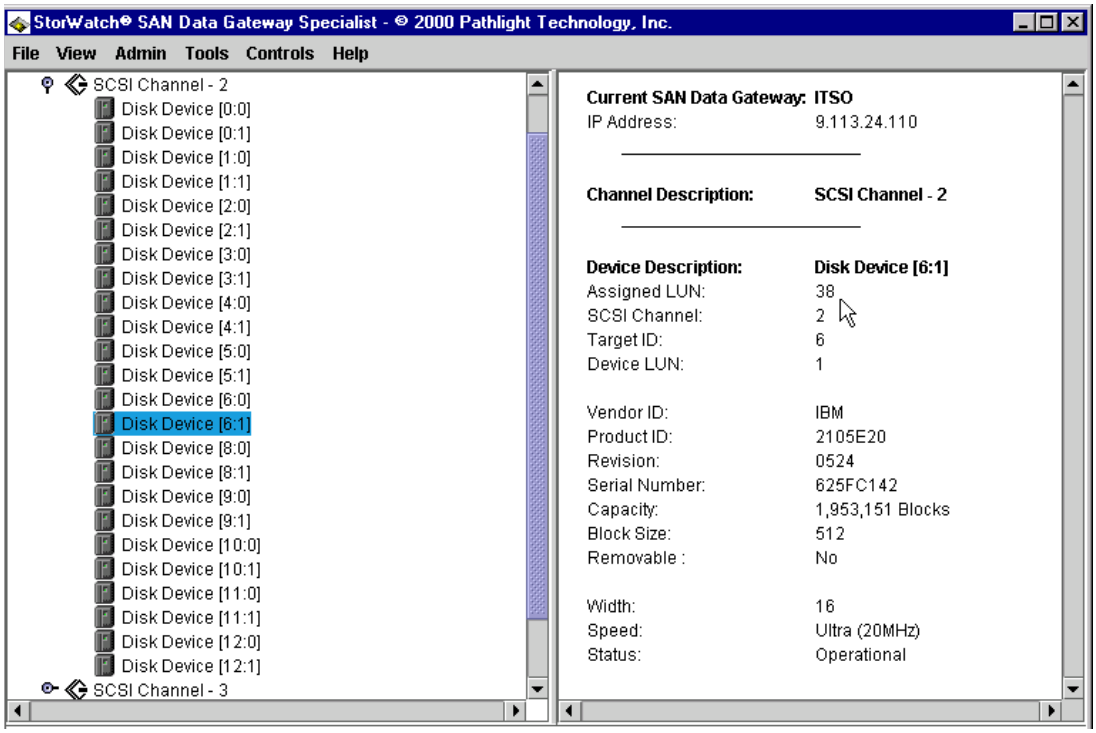

Figure 236. SCSI LUN assignment

In this example, the disk device attached has a SCSI target of six and LUN ID of 1. When it was recognized by the SAN Data Gateway, the Gateway assigned LUN number 38 to this device. This Gateway's assigned LUN number corresponds to the assigned LUN number that appears on the top of the VPS settings window as shown Figure 238 on page 282.

8.5.2.6 SCSI to LUN map

Prior to setting the LUN masking, it makes sense to prepare a list that has each LUN and shows which host is allocated to that LUN. To assist in this process the Gateway has a map that provides a cross reference of the actual SCSI target and ID to the Gateway's assigned LUN number. To access this map you must access the service terminal from the serial port as described earlier in 8.2, "Installation" on page 252. Once connected, type in the command `mapShowDatabase`. The database lists not only devices that are presently connected, but also devices that have previously been connected. If a previously attached device is later reattached, it is assigned back to its

previous address. Figure 237 shows the output returned to the service terminal.

```
ITSO > mapShowDatabase
devId  Type  Chan  tId  tLun  UID
-----------------------------------------------
000    SNA   127   127  127   00000060:45160e08
001    SCSI  001   000  000   20100060:45160e08
002    SCSI  001   000  001   20200060:45160e08
003    SCSI  001   001  000   20300060:45160e08
004    SCSI  001   001  001   20400060:45160e08
005    SCSI  001   002  000   20500060:45160e08
006    SCSI  001   002  001   20600060:45160e08
007    SCSI  001   003  000   20700060:45160e08
008    SCSI  001   003  001   20800060:45160e08
009    SCSI  001   004  000   20900060:45160e08
010    SCSI  001   004  001   20a00060:45160e08
011    SCSI  001   005  000   20b00060:45160e08
012    SCSI  001   005  001   20c00060:45160e08
013    SCSI  001   006  000   20d00060:45160e08
014    SCSI  001   006  001   20e00060:45160e08
015    SCSI  001   008  000   20f00060:45160e08
016    SCSI  001   008  001   21000060:45160e08
017    SCSI  001   009  000   21100060:45160e08
018    SCSI  001   009  001   21200060:45160e08
019    SCSI  001   010  000   21300060:45160e08
020    SCSI  001   010  001   21400060:45160e08
021    SCSI  001   011  000   21500060:45160e08
022    SCSI  001   011  001   21600060:45160e08
023    SCSI  001   012  000   21700060:45160e08
024    SCSI  001   012  001   21800060:45160e08
025    SCSI  002   000  000   21900060:45160e08
026    SCSI  002   000  001   21a00060:45160e08
027    SCSI  002   001  000   21b00060:45160e08
028    SCSI  002   001  001   21c00060:45160e08
029    SCSI  002   002  000   21d00060:45160e08
030    SCSI  002   002  001   21e00060:45160e08
```

Figure 237. Service terminal display of device map

The numbers on the left are the assigned LUN numbers, and note that number ‘0’ has been assigned to the Gateway. The other columns contain the device type, the SCSI channel it is connected to on the SAN Data Gateway, and the actual target and LUN ID.

8.5.2.7 Setting host access to LUNs

To view and set host access to particular LUNs, access the Virtual Private SAN Access Settings by selecting **Controls -> Access Options -> Virtual Private SAN**. This window will show all the hosts that have registered to the Gateway. To allow a host to access a particular LUN, place a check mark in the row that corresponds to the host. To disable access, the square must be clear, without a check mark. The check mark is toggled on and off by clicking in each square. Figure 238 provides an example of LUN masking.

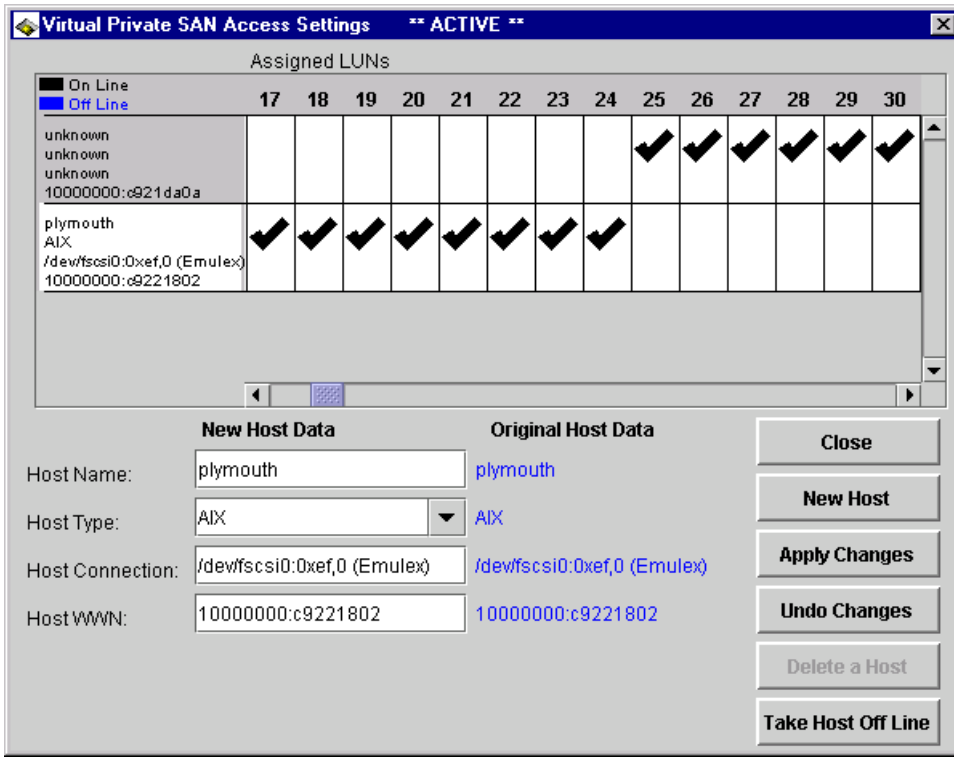

Figure 238. Setting LUN masking

The host *plymouth* has access to LUNs 17 to 24 but does not have access to LUNs 25 to 30. Conversely, the host at the top has access to LUNs 25 to 30, but does not have access to LUNs 17 to 24. Using the scroll bar on the bottom, we can scroll through all 255 available LUNs in the SAN Data Gateway, and enable or disable access by adding or clearing the check mark. The scroll bar on the right hand side allows us to scroll through the different hosts. In this example, there are only two hosts, but there can be several more.

Once completed, select **Apply Changes** and then **Close**. The host system may now have to be restarted or some other method used to rescan the bus for the host to detect that it now has access to new devices.

Any combination is allowed, so if the same LUN is to be shared by two different hosts, a check mark for that LUN must be set for both. If this is the case, the host systems must have a device sharing software installed to control access to the disk device for data integrity.

As each Fibre Channel port can support up to 8 eight different hosts, there can be up to 48 hosts attached to the SAN Data Gateway. They can all share the same four SCSI channels. By using the Virtual Private SAN feature, you can ensure that only the LUNs you desire a certain host to use will be accessed and that no other host will access them.

8.5.3 Combining Zoning and Virtual Private SAN

If Virtual Private SAN is enabled and LUNs have been allocated to specific hosts, then zoning is not necessary or required. The Channel Zoning window can remain at the default settings with all Fibre Channel ports accessing all SCSI channels.

However, they can work in combination to add an extra level of control and security. If zoning is added, then VPS can only control the LUNs that are included in its access zone. Figure 239 shows a zoned SAN Data Gateway.

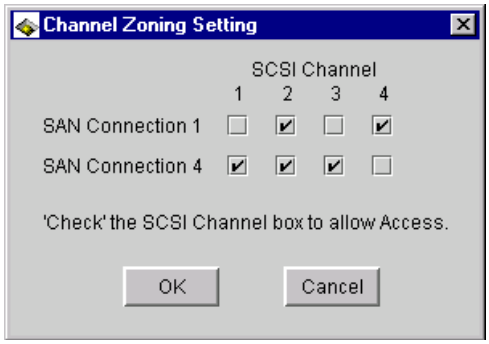

Figure 239. Combining channel zoning and VPS

SAN connection 1 has access to the LUNs on SCSI channels 2 and 4. VPS will control access to LUNs on SCSI channels 2 and 4 for SAN connection 1. Since SAN 1 is *not* zoned for SCSI channel 1 and 3, it will not reach any LUNs on these channels. Even, if a LUN on SCSI 3 has been enabled for access to a host on SAN 1 in the VPS settings window, the host will not see that LUN because of the zoned settings. The same is true for SAN connection 4. A host connected here will access LUNs only on SCSI channels 1, 2, and 3, but not channel 4, regardless of whether the VPS settings will allow it.

8.6 Adding Fibre Channel fabric components

The SAN Data Gateway is an integral component in a storage network. As such, you can also attach other SAN fabric components to the Gateway to

increase the connectivity options in a SAN. Hubs and switches can be easily added, and allow many more systems, local or remote, to access the SCSI devices on the Gateway.

As switches provide more flexibility, and hubs are mainly used to extend distances, the following discussion will focus on switch implementation, rather than hubs.

8.6.1 Connecting an IBM SAN Fibre Channel Switch

The IBM SAN Fibre Channel Switch provides either an 8 port or 16 port switch that can also be cascaded.

8.6.1.1 Allowing fabric connection

Before connecting a switch to the SAN Data Gateway, there is a setting that must be checked to allow proper fabric connection. From the SAN Data Gateway Specialist, select the Fibre Channel port that will have the switch connected. Select **Control -> Fibre Channel**, and the Set Fibre Channel Parameters window appears, as shown in Figure 240.

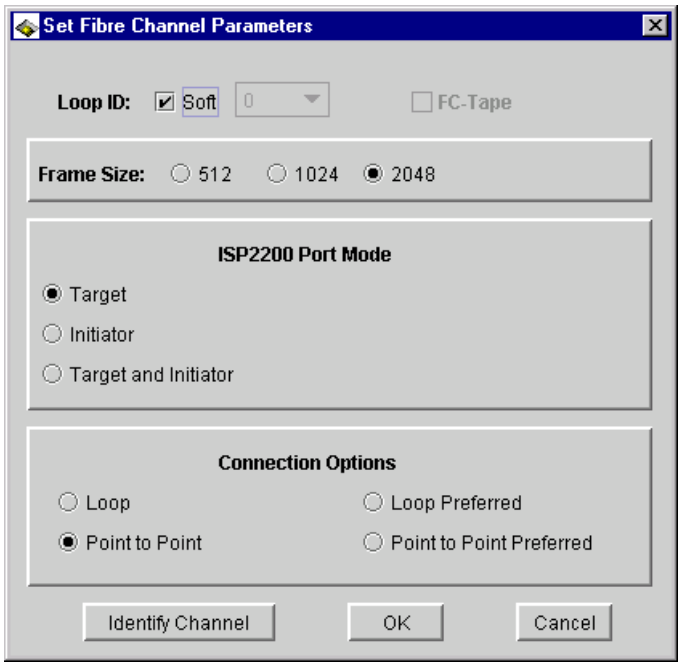

Figure 240. Fibre Channel port setting for switch attachment

By default, the setting in the **Connection Options** box will have **Loop** enabled. For switch or fabric connection, select **Point to Point**, and then

select **OK**. The SAN Data Gateway must be restarted for the change to take effect.

Note

If a **Connection Options** box does not appear, the Fibre Channel module installed will support loop connection only and will not support fabric connection. The module must be replaced or select a port that will support fabric connection.

A Fibre Channel cable can be connected from the switch to the port on the Gateway. The data to the Specialist must be refreshed by selecting and highlighting the Gateway and then select **View -> Refresh SAN Data Gateway.**

8.6.1.2 Switch registration

With VPS enabled, the switch will register with the database the WWPN of the port on the IBM SAN Fibre Channel Switch. Figure 241 shows a switch connected to port 4 of the Gateway, and which has registered to the database. Since we cannot load any host registration software onto the switch, all other fields are left *unknown.*

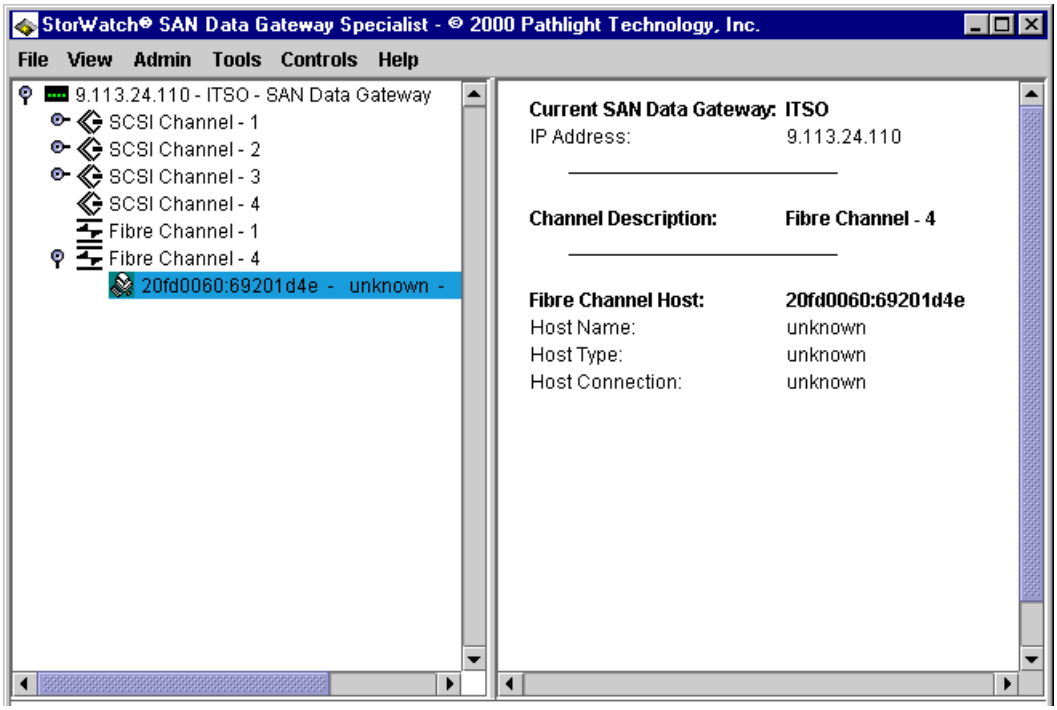

Figure 241. Switch registration

By selecting and highlighting the Fibre Channel port, as shown in Figure 242, you will see in the right hand view pane that the port is now using a point to point, or N_port, connection that denotes a fabric login rather than a loop login.

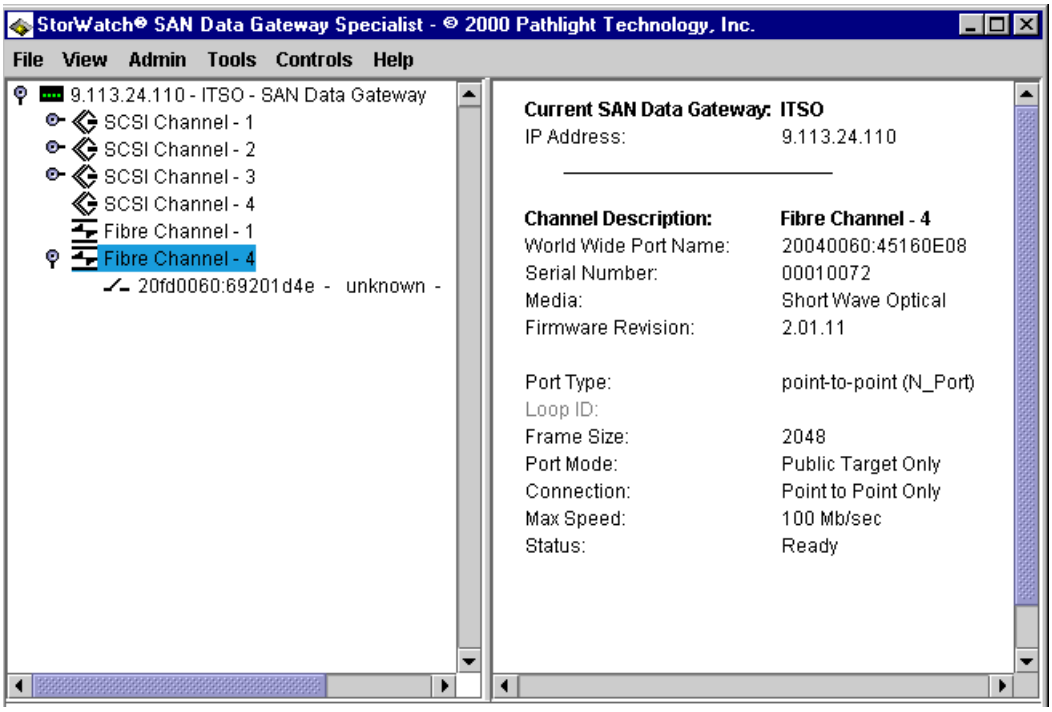

Figure 242. Switch port login

You can also check the port on the switch, by accessing the switch configuration from a Web browser, to ensure that the port has registered and is communicating properly. Figure 243 shows that, in this example, port 7 of the switch was used to connect to the Gateway. Note that the WWPN is the same as in the Specialist window and that the port type is F_port. For more information on the switch, refer to Chapter 6, “Implementing an IBM SAN Fibre Channel Switch” on page 139.

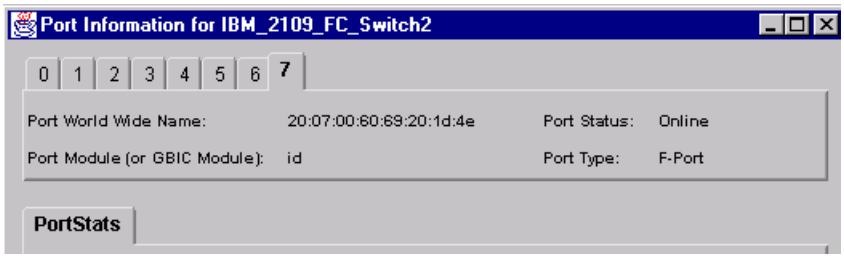

Figure 243. IBM SAN Fibre Channel Switch port settings

8.6.1.3 Changing the switch information

From the VPS Access Settings window, it is possible to change the *unknown* information of the switch. **Select Controls -> Virtual Private SAN**, and select the entry that has the WWN of the switch. You can now change the information to further describe the switch, or other pertinent information if desired. This is shown in Figure 244.

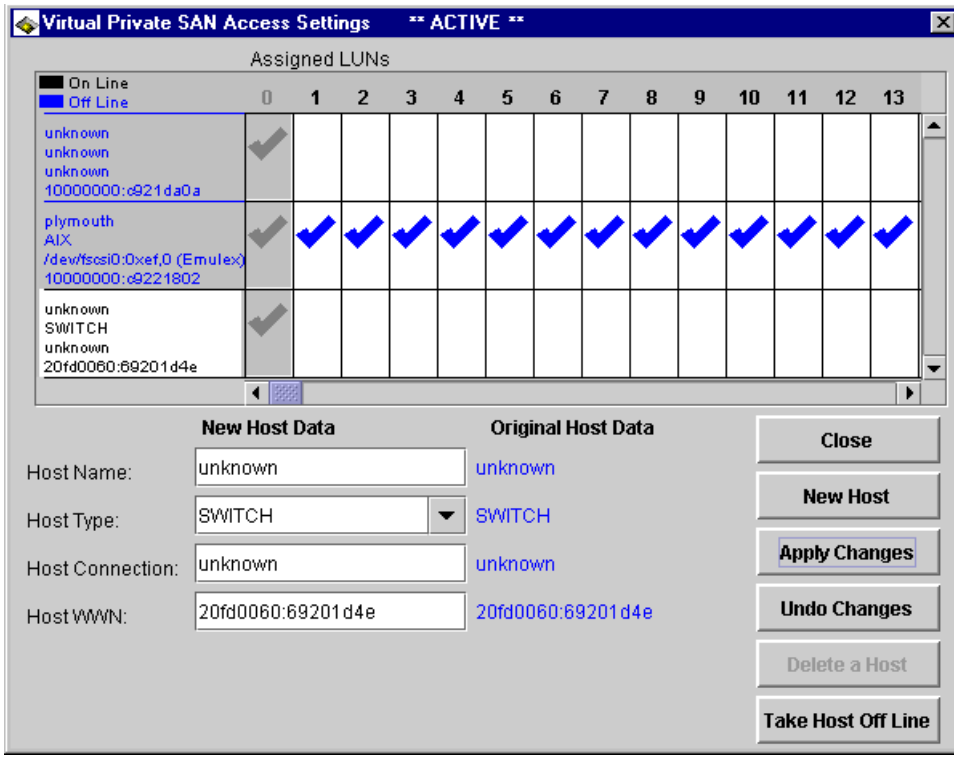

Figure 244. Changing switch information

The information on the host *plymouth*, that was attached using port 4, is still kept in the VPS database, but it is now shown in blue to indicate that it is offline.

Figure 245 is a diagram to show the configuration with the switch.

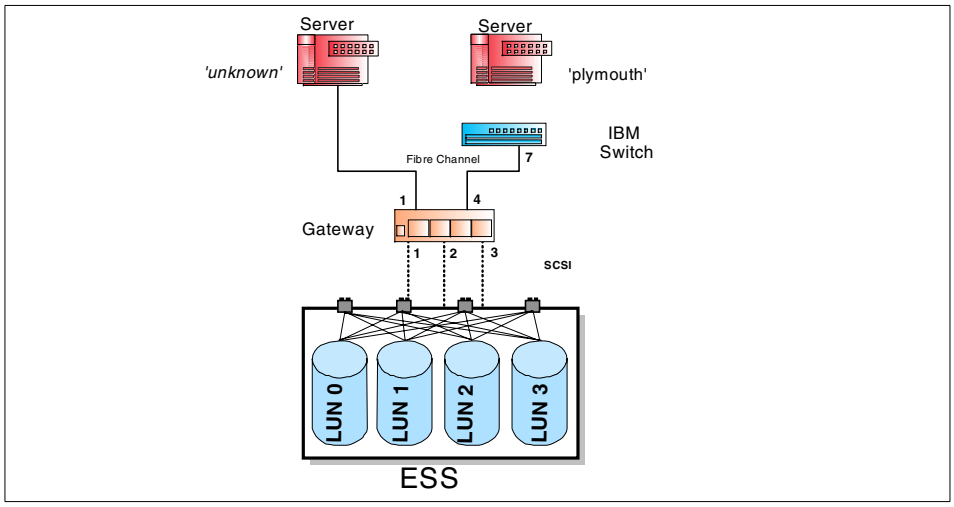

Figure 245. Adding an IBM switch to the Gateway

Once changed, as shown in Figure 246, the information shown on the main window will reflect the change, and the icon on the left hand side of the WWN changes to depict a switch.

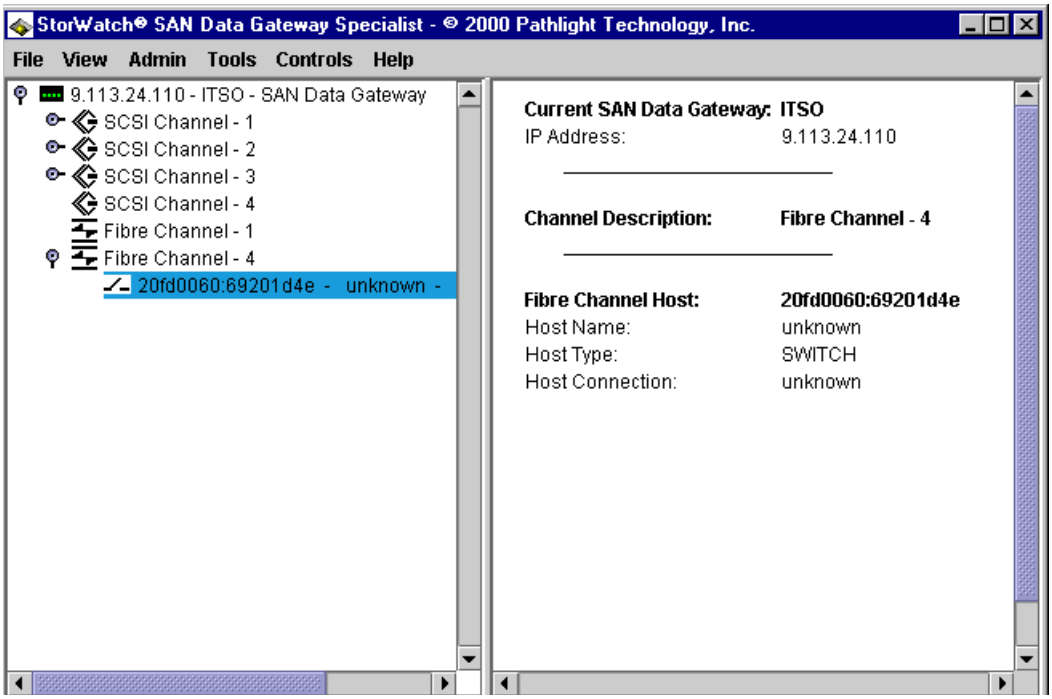

Figure 246. Switch port information

8.6.1.4 Attaching hosts to the switch

Any hosts that will attach to the switch should have the host registration software loaded and installed. Refer to 8.5.2.2, “Host registration” on page 273 for details. Plug in the Fibre Channel cable from their respective adapters and power on, or restart the host. The registration of the hosts will be completed through the switch. Figure 247 shows a switch installed on port 4 and two hosts connected on the switch. The host *plymouth* is now connected through the switch and an NT host was added.

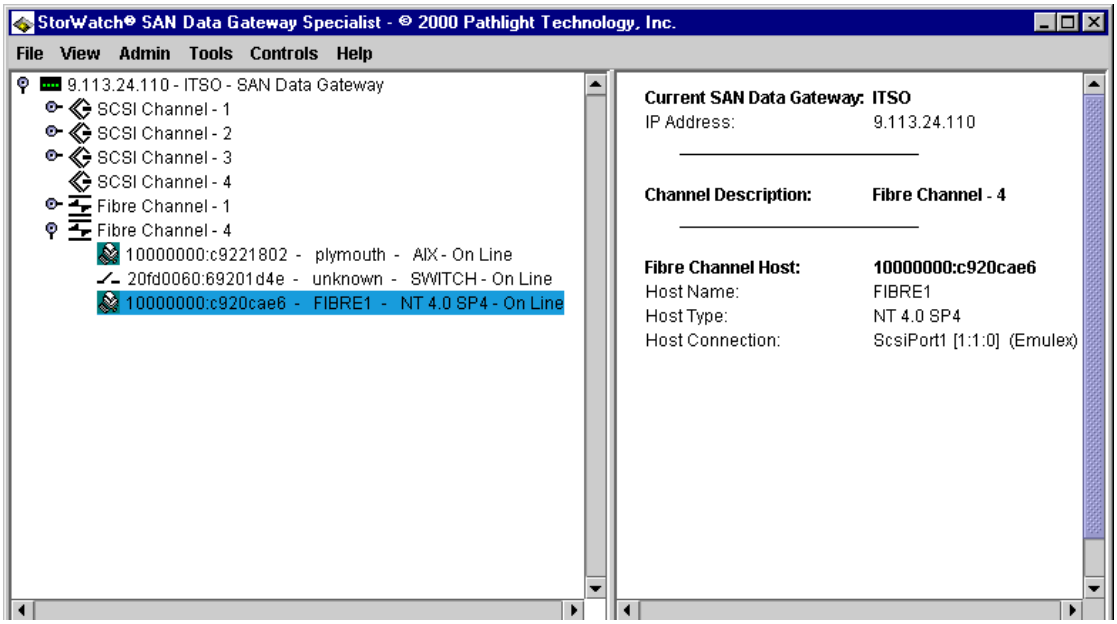

Figure 247. Hosts and a switch on a Gateway port

By selecting and highlighting the new NT host, we can see its information, that was automatically sent by the host registration process.

The configuration with two different hosts connected to the switch, as described previously, is shown in Figure 248.

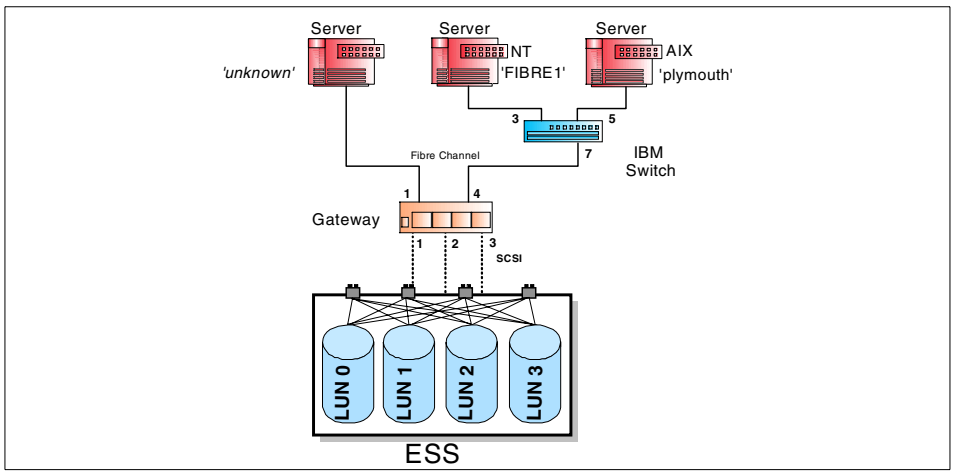

Figure 248. Adding two heterogeneous hosts to the switch

8.6.1.5 Device access

To have the new NT host access some of the devices, you need to set up the VPS Access parameters by selecting **Controls -> Access Options -> Virtual Private SAN**. Figure 249 shows the switch and the two hosts in the VPS database. As *plymouth* was previously connected direct to the Gateway, its settings have been maintained, but now that it is reconnected, it is back online. The NT host does not have any LUN access yet.

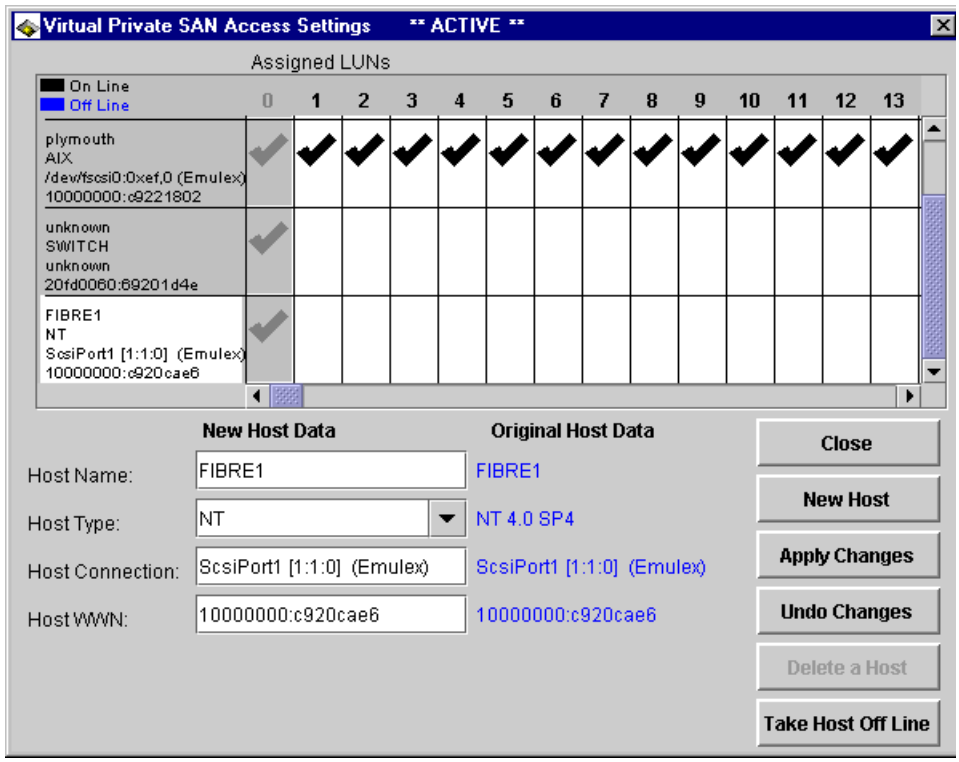

Figure 249. VPS Access window with switch and two hosts

A check mark in the box allows a host access to a LUN, or clear the box if you want to restrict access as described in 8.5.2.5, “Setting up a Virtual Private SAN” on page 279. Figure 250 shows the host *plymouth* with access to assigned LUNS 17 to 24, and the NT host *FIBRE1* is now set with access to LUNs 25 to 30.

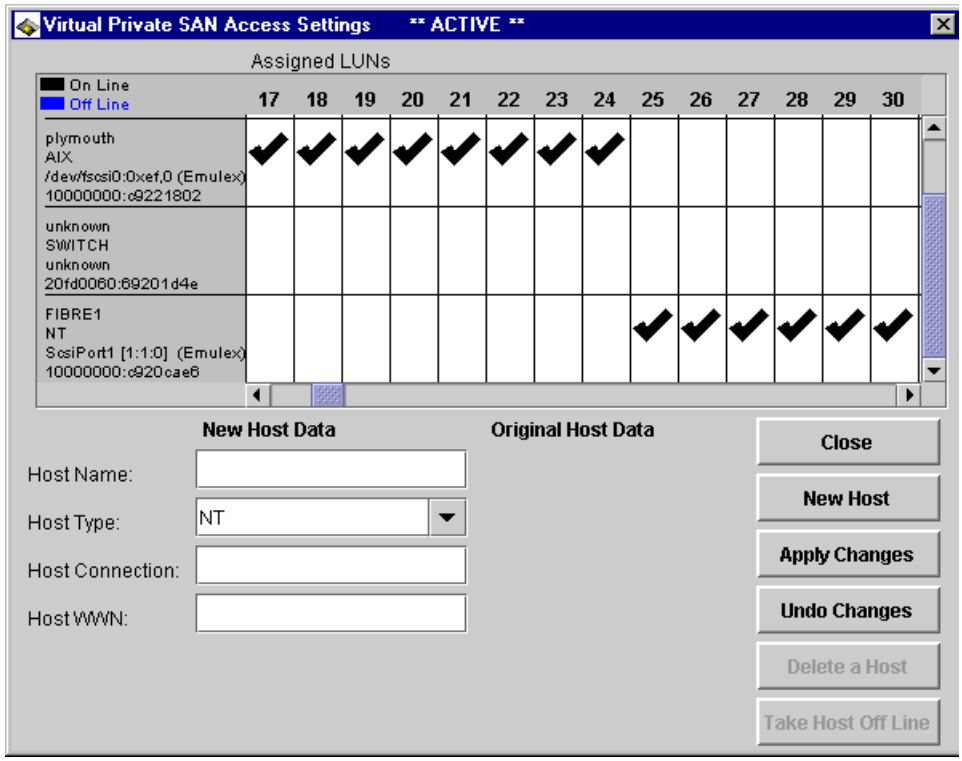

Figure 250. Setting LUN access for the host FIBRE1

Once you select **Apply Changes** and then **Close**, the new settings will be in effect. For Windows NT to recognize the new devices, it will have to be restarted, so that it will do a rescan.

The 'SWITCH' entry in the VPS database does not require any check marks to be set. Because the WWN of the host is known, and the switch WWPN is known, LUN access is accomplished through the switch as if it was transparent.

Note

LUN access is not only affected by the VPS and Zoning with the SAN Data Gateway, there is also Zoning within the IBM SAN Fibre Channel Switch. Be aware of the zone settings within the IBM SAN Fibre Channel Switch.

8.6.1.6 ESS considerations

As you configure and allow access from host systems to volumes that are in the ESS through the Gateway, you must consider how the volumes were created within the ESS.

Volumes in the ESS are assigned to a particular SCSI port in the ESS. As they are assigned, you also specify the host type that will access these volumes.

Through the SAN Data Gateway it is possible to assign volumes to a specific operating system in the ESS, but then have a host with a different operating system access to these volumes.

As an example, look at Figure 248 on page 291. There we have a Windows NT and an AIX host. The ESS was configured and the Gateway attached so that SCSI 2 was assigned Windows NT volumes and SCSI 3 was assigned AIX volumes. It would be possible in the SAN Data Gateway to allow the NT host to access a LUN on SCSI 3 and the AIX host access to volumes on SCSI 2.

Here is where the device map would again be helpful in determining what Gateway assigned LUNs should be assigned to each host.

Volumes within the ESS assigned to a particular SCSI port should be of the same host type.

8.6.2 Connecting a McDATA Enterprise Fibre Channel Director

The McDATA Director can be used to attach to the SAN Data Gateway. The connection of the McDATA Director to the SAN Data Gateway is similar to the description provided previously in 8.6.1, “Connecting an IBM SAN Fibre Channel Switch” on page 284.

The difference that can be seen when using the McDATA Director is, that once connected, the Director does not register the WWPN of the port connected to the Gateway *visually* to the VPS database. Figure 251 provides an example of where a McDATA Director was connected to the Gateway on Fibre Channel port 4.

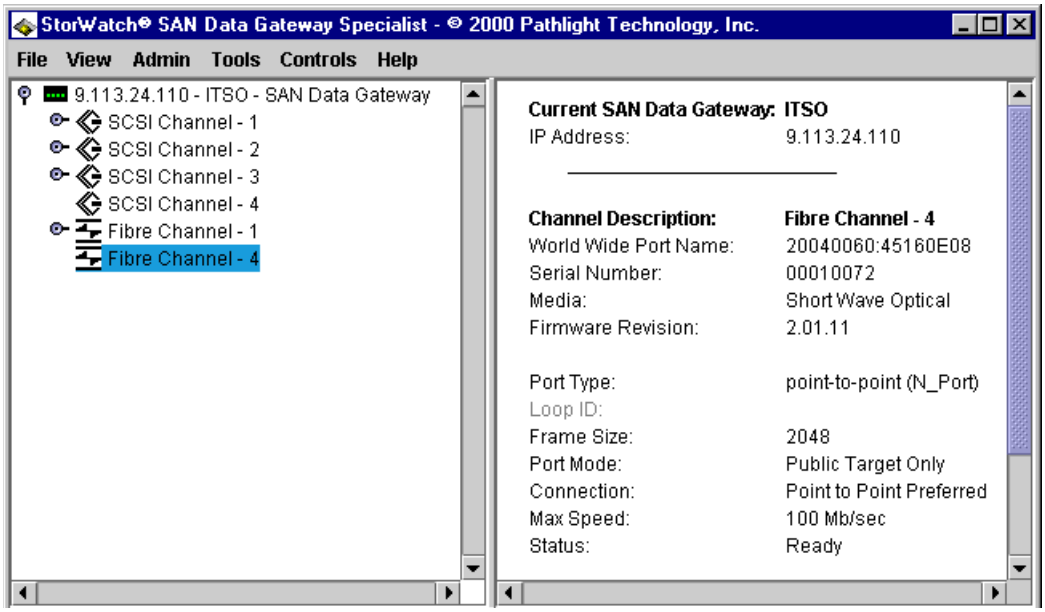

Figure 251. McDATA Director connection to a Gateway

In the right hand view pane under Port Type, you see that the port is in N_Port mode to denote a fabric connection. However, in the left hand view pane there are no devices listed under the Fibre Channel port 4 description.

As hosts begin to attach to the McDATA Director and are restarted, they will begin to login their own information to the Gateway's VPS database. The McDATA is seen as completely transparent to the SAN Data Gateway.

8.7 High availability considerations

Connections from the host to the SAN Data Gateway and from the Gateway to the ESS, or other SCSI devices, have been discussed until now with single host adapters or utilizing single SCSI connections. Special considerations must be taken into account in deciding to add host adapters, or providing redundant SCSI connections.

8.7.1 Single host with multiple Fibre Channel adapters

A host can have several adapters installed. If each adapter is assigned to different LUNs using the SAN Data Gateway, there are no contention issues. As each adapter is assigned specific LUNs, it cannot see or access any other LUNs that may be present.

However, there may be situations where the adapters are used to perform load balancing, and failover much like the SCSI connections from a host to the ESS that are in place today. As the host in the SCSI situation requires the Subsystem Device Driver (SDD) on each host, this is also true when using the Fibre Channel adapters.

Software, similar to SDD, is required in the host for it to recognize that it may have two or more paths to the same volume. If this was not loaded on the host, the host would recognize the volumes as being different, and there will be contention problems, and data corruption.

As SDD is currently only supported on SCSI adapters and not with Fibre Channel adapters, it is also not supported when using the SAN Data Gateway to connect to an ESS.

8.7.2 Multiple SCSI connections

The SAN Data Gateway can have all four SCSI channels connected to a single ESS. A volume in the ESS can be assigned to more than one SCSI port in the ESS. If these SCSI ports, that have the same volume assigned to it are connected to a SAN Data Gateway, the Gateway will assign multiple LUN ID numbers to the same volume. This is because the Gateway, upon discovery, or scanning of the SCSI buses, will view the volumes on each SCSI channel as separate volumes. For further explanation on this refer to 8.4.2, “Device discovery” on page 270.

If each LUN ID was then masked and zoned to different host systems, it is vital that the hosts have some access sharing software loaded to control access and avoid data corruption.

If the LUN IDs were assigned to the same host, then again software similar to SDD is required for the host to recognize that it has multiple paths to the same volume.

As stated earlier, SDD is not supported on the SAN Data Gateway connection to an ESS.

8.7.3 Adding Fibre Channel switches

As switches are added to allow for more host attachments, there are considerations about how many paths the host has to the SAN Data Gateway. If we refer back to Figure 248 on page 291, there is a single path from the switch to the Gateway.

If another Fibre Channel path from the switch to the Gateway was added, each host now has two paths to access the same LUNs. Each host will *see* the volumes twice. Once again, to prevent the host from recognizing the same volume twice, software similar to SDD is required.

Another option here is to utilize the zoning and LUN masking capabilities of the SAN Data Gateway. This would ensure that certain volumes can only be accessed on one Fibre Channel Gateway port and by a particular host.

Also available is to add zoning within the switch. The switch Fibre Channel ports can be zoned so that the host only has one path to the SAN Data Gateway. This would be used in combination with the zoning and LUN masking features of the SAN Data Gateway.

Chapter 9. Implementing the Vicom Fibre Channel SLIC Router

The Vicom Fibre Channel SLIC Router, 7139-111, enables all IBM 7133, 7131, and 3527 SSA Serial Disk Systems to attach to host systems using Fibre Channel host adapters and drivers. This allows you to protect your investment in SSA disk, while being able to create and build a SAN infrastructure.

The Vicom Fibre Channel SLIC Router replicates data across or within serial disk systems — simultaneously mirroring two or three copies of data without host involvement. With global hot disk sparing, data is automatically rebuilt if a mirrored disk fails. In this way, the Vicom Fibre Channel SLIC Router improves performance and data availability while simplifying storage operations.

The Instant Copy function can create a separately addressable copy of mirrored data that can be used for tape backup. After the backup has completed, data is resynchronized with the primary copy.

The Vicom Fibre Channel SLIC Router also can create composite drives by concatenating up to 16 physical disks.

Using these functions, physical drives become members of larger or more complex logical drives.

A diagram to depict a single host to router configuration is shown below in Figure 252.

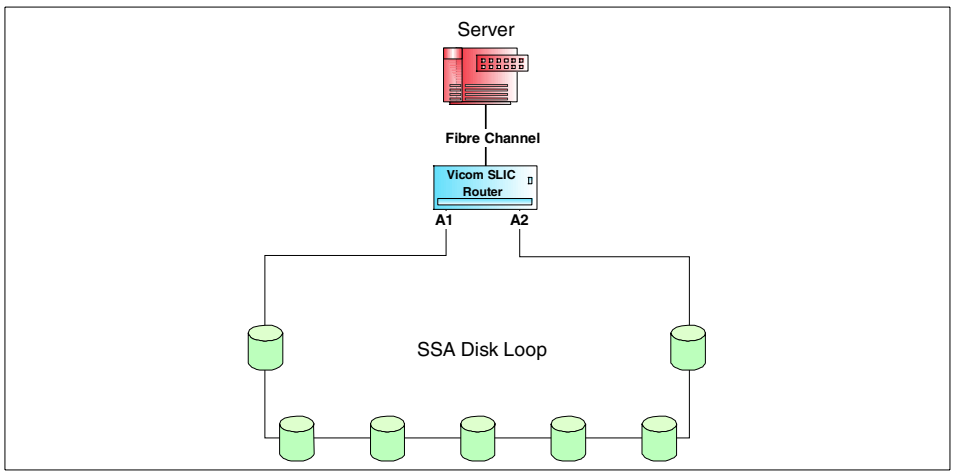

Figure 252. SLIC Router with a single host

In the following sections, we cover these topics:

9.1 Installing the SLIC Router

To install and establish a SLIC storage system, the Router and all the disks to be used must be setup in a proper sequence. During this sequence, only *a single* SLIC Router must be used to do the configuration. Once configured, other SLIC Routers can be added to the loop. This sequence is described below.

Note

For information and a description to understand the LED codes that will be discussed, please refer to the *SLIC Router Installation and User's Guide*, 310-605759.

1. Power on SSA disks

   The SSA disks must be powered on to ensure that all disks spin up and pass the power sequence. Make sure all drive LEDs are on solid to indicate a functioning device. Any faulty or suspect disk drive should be replaced.

2. Clear Router Node Mapping

   To begin the Router installation, the first step is to clear the Node Mapping table. This is done by *shorting* the SSA ports on the rear of the router. Plug in an SSA cable from port A1 to A2 on the back of the Router and power it on. Clearing the table will take only seconds, and when completed, the Status LED on the front of the Router will flash a code 060. The router is now powered off, the SSA shorting cable is removed and the SSA disks are attached.

3. Connect SSA disks to Router

   All the disks to be used are to be connected together to form a complete SSA loop with the Router included. All dip switches in SW1 should be in the down position. On SW2, dip switches 0 and 1 are set in the down position, all other switches in dip switch 2 should be turned up. This is

considered *mode 3*. Power on the Router, the Status LED will begin to flash rapidly as it searches the SSA loop to recognize all the disk drives. This may take approximately 1 minute to complete. Once the Status LED has stopped flashing and is solid, the process is complete and the Router is powered down.

4. Router Subsystem Diagnostic test

   A Subsystem Diagnostic test is now run on the disk drives called *mode 15*. This tests the disk drives for spindle spin up, read tests and nondestructive write tests. The Router is set for mode 15 by setting switches 0, 1, 2, and 3, on SW2, to the down position, and the rest turned up. The Router is now powered on, the Status LED will flash rapidly. The test will be done on each disk drive in the SSA loop separately and will begin with the drive closest to the A1 port on the back of the Router. As the test is completed on a drive, the LED on the SSA drive will flash and then it will move to the next drive. This test should continue until all drives have been tested. The test runs in a continuous cycle, so once all drives have been tested at least once, the Router is powered off.

   If a drive fails the test, the testing will stop, and the Router's Status LED will flash a diagnostic code. A code map with a description of the errors can be found in the *SLIC Router Installation and User's Guide*, 310-605759.

5. Assign Fibre Channel target

   With the Router powered off, you can now assign a Fibre Channel target ID number to the Router. Any number can be selected, however, this number must be a unique ID. No other device can have the same Fibre Channel target ID once it is set on the Router.

   This is done by setting selected dip switches in SW1 to the down position. The switch is set up in binary notation: a switch that is down represents a 1 and a switch up represents a 0. Figure 253 shows the switch numbers and their corresponding value.

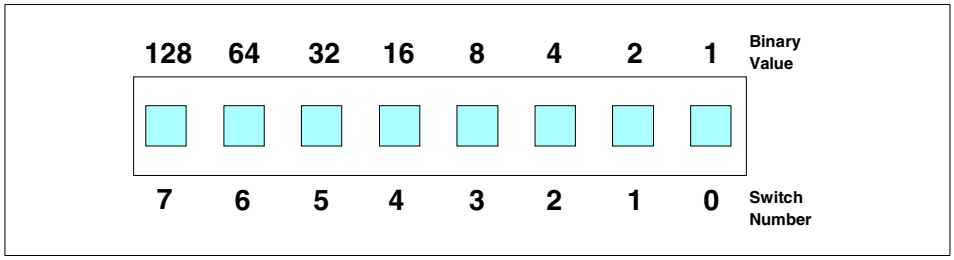

Figure 253. SW1 dip switches

For example, by setting switch 2 and 0 down, a value of 5 is represented. All other switches would be in the up position. By setting a switch down, its value is added. To change a number that was previously set on a Router, power must be cycled to the Router for the change to take effect.

6. Mapping physical drives

   Before powering on again, SW2 must be set to mode 3 with switch 0 and 1 set down, and all other switches set up. The Router is powered on, the Status LED will flash rapidly to rediscover the disks and the SSA drive LEDs should be on solid. Once completed, the status LED will be solid, and the drives are now considered to be mapped physical drives. The Router is powered off.

7. Create general spares

   The mode on SW2 is changed to mode 12 to set all drives to be general spares. Mode 12 is represented by setting switch 2 and 3 down and the rest turned up. Power on the Router again, the Status LED will flash. After approximately one minute, the LED will flash code `100` to indicate it has completed. The Router is again powered off.

8. Format the drives

   The Router is set to mode 14, switch 1, 2, and 3 down on SW2, to format all disk drives. Power on the Router, depending on the number of drives, the format process will take anywhere from 30 to 60 minutes. During this time, the Status LED will flash rapidly and the SSA drive LEDs will flash. When it has completed, the Status LED will flash code `100`. Power off the Router.

9. Clear node map

   Once completed, the node map *must* be cleared. This is done as described earlier by *shorting* ports A1 and A2 with an SSA cable. Power on, wait for code `060` to flash and then power off.

   The drives are now ready to be assigned and used on a host.

   You can also now setup mirror drives or composite drives within the Router. This is done by setting the switches in SW2 to other modes. For detailed information on setting the switches and selecting the other modes, please refer to the *SLIC Router Installation and Users Guide*, 310-605759.

10. Host attach and power up sequence

    For a host to now recognize and use the disks, set the dip switches in SW2 back to mode 3, this is normal host operation mode. The Fibre Channel cable from the host can be connected to the Router. If the SSA drives are

not powered on, do this now, and this should be done before the Router. Next, the Router is powered on, wait for the Status LED to stop flashing and remain on solid. At this point the host can be powered on.

A check can be done to see that the SLIC Router is being recognized by the host adapter card. On a Windows NT with a QLogic Fibre Channel adapter, during boot up look for a prompt to enter the QLogic bios by entering in ALT Q. At the bios window, select **Scan Fibre Devices**. A list of the Fibre Channel target IDs are presented, scroll down to the ID that you set in SW1. You will see the WWN of the SLIC Router. Exit the bios and the system will reboot.

Once the system has started, you use a method to ensure that the host has access to all the drives. This is different depending on the operating system of the computer. For Windows NT, select **Start -> Programs -> Administrative Tools -> Disk Administrator**. This tool will report that new disks have been found and will be added to the system.

9.2 SLIC Manager software

Rather than using the dip switches to configure the features, another option is to use the SLIC Manager software. The SLIC Manager also provides configuration, monitoring and management capabilities of the SLIC router and the SSA drive loop. The SLIC Manager can be setup to allow remote access if desired.

The Manager software consists of a server and client portions. The server includes a daemon service and a user interface. The client has the user interface only.

The server portion must be loaded on the host that is directly attached to the Router, as the daemon service is started from here. The daemon must reside on the host that is directly connected to the Router. This host can also be used to run the Manager software for local access.

The client software can be loaded on to any computer, running a supported operating system, that can communicate to the host with the daemon service running. It must communicate to the server host using TCP/IP. This allows remote access to the Router and the storage loop. This is depicted in Figure 254.

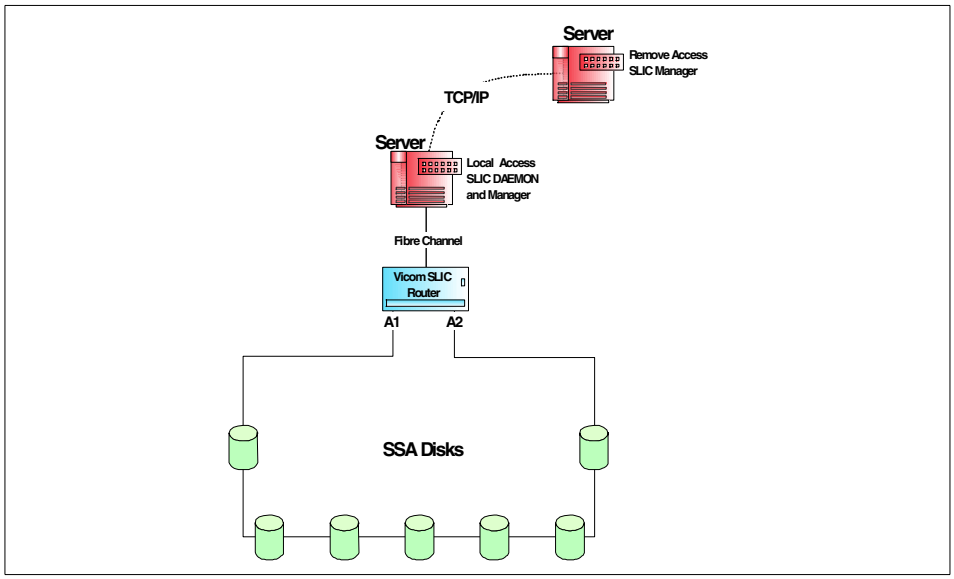

Figure 254. SLIC Manager access

The SLIC Manager has a GUI and a command line interface available for Windows NT systems. UNIX systems will only have the command line interface available.

The following installation and configuration examples will detail using the GUI from a Windows NT platform. To review the commands that are available for UNIX platforms and installation instructions on other operating systems, please refer to the *SLIC Manager Installation and User Guide*, 310-605807.

9.2.1 Installing the SLIC Manager software

The SLIC Manager software can run on many operating systems. The following discussion will describe an installation on a Netfinity 5500 with Windows NT 4.0.

To install the SLIC Manager server software for local access, the Vicom Utilities CD-ROM is placed in the CD drive. Select **Start -> Run** and **Browse** the CD drive. Go to `slicmgr\i386\server\setup.exe` and click **OK**. Follow the prompts displayed on the window to install the Manager software. This will install the daemon service also.

For remote or client access the daemon service is not required. To load the Manager software only, go to `slicmgr\i386\client\setup.exe` instead.

9.2.2 Communicating to the Router

For the SLIC Manager server software to communicate to the Router, it requires space on a disk or several disks that are within the SSA loop. This is referred to as a SLIC Zone. To create space on a disk, a file or partition — depending on the operating system used — is created for the Manager software to use. To create this SLIC Zone, a configuration file must be created or edited.

9.2.2.1 Editing the configuration file

When the software is loaded, a sample configuration file called *7190.cfg* is added in the `C:\ibm7190\sdus` directory. This is a text file that can be viewed and edited by simple text editors, such as Windows Wordpad. Open up the 7190.cfg file and it will contain a sample of how the file should look. Also note that on the left hand side, the '#' sign is entered in every line to mark it out as a comment. This is shown in figure Figure 255.

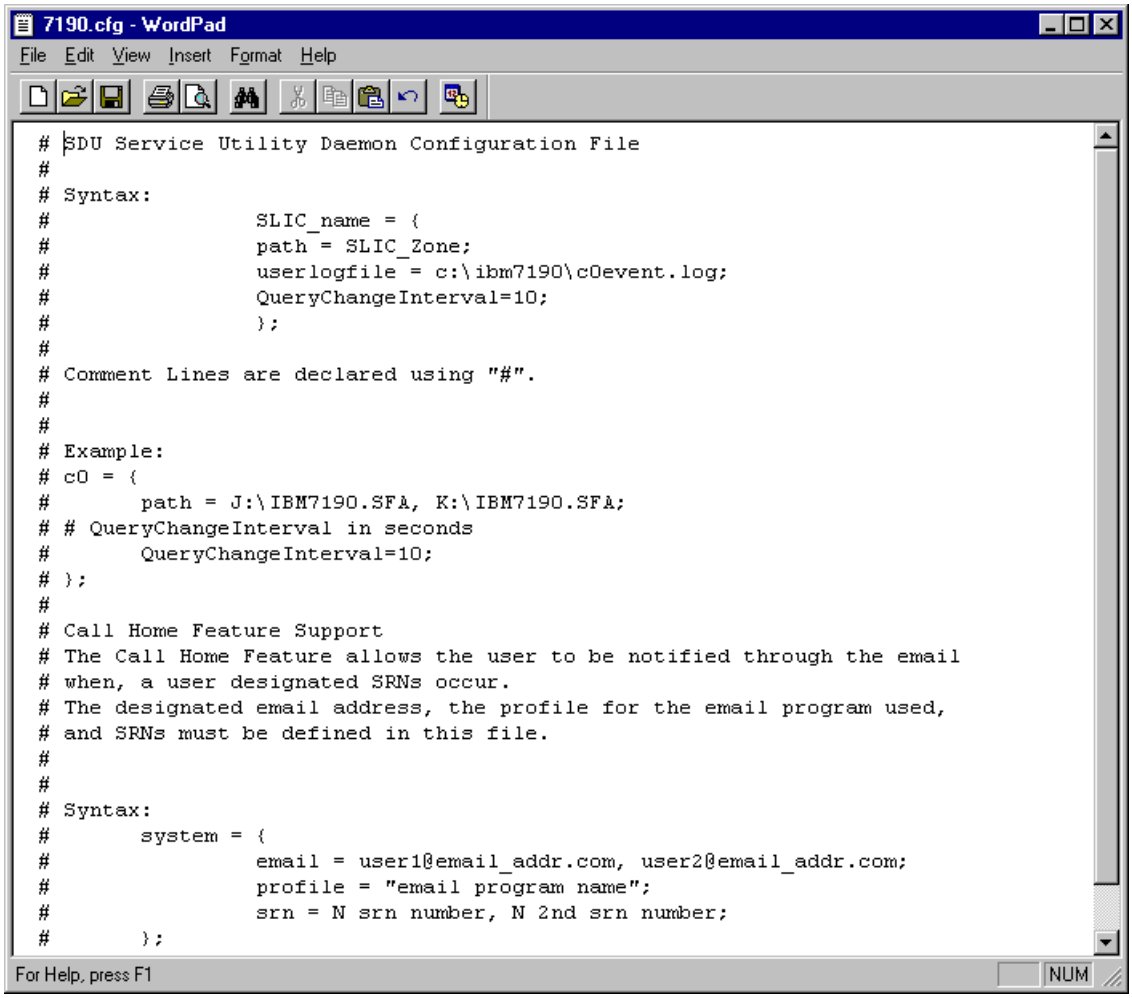

Figure 255. Sample configuration file

This file can now be edited to be used as the configuration file for your SLIC system. Begin by deleting the '#' sign on the lines that contain the sample configuration. The rest of the information can be entered in, as shown in Figure 256. A description of each entry field is also provided.

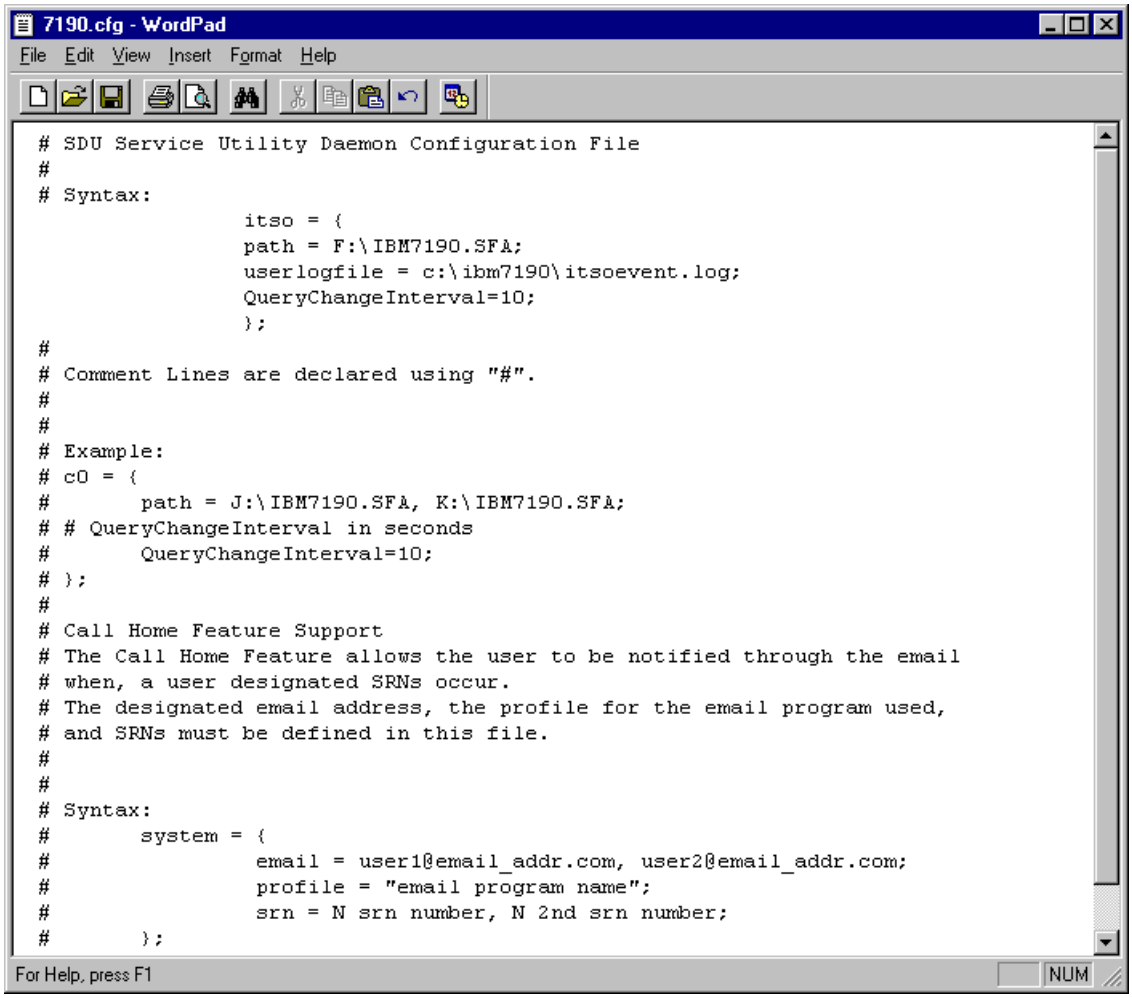

Figure 256. Edited configuration file

9.2.2.2 Configuration file information

The *SLIC_name* can be any name that you would like to use to identify the Router.

Creating a SLIC Zone

The *path* refers to the SLIC Zone, file or partition, used for the Manager to communicate to the Router. To edit this option, it requires that a drive on the SSA loop has been recognized by the host and that the drive has been formatted. In the example above, a Windows NT host was used. The

Windows NT Disk Administrator was used for the host to recognize the drives, the first drive assigned the next drive letter, F, and it was formatted.

The file naming for a SLIC Zone depends on the type of operating system running. For Windows NT, the naming is <drive letter>:\IBM7190.SFA. You can enter in many SLIC Zones, but only one is required to get access at the beginning. After the other drives have been configured as mirrors or composite drives, then SLIC Zones can be created for these drives if desired.

Including many SLIC zones in the path statement will allow the Manager to access a zone on another drive. This is helpful to protect against when a drive fails, and that drive has a SLIC zone defined to it. If the Manager cannot access the first SLIC zone, it would try the next zone in the order it was entered in the path statement.

For the naming conventions used on other operating systems to create a SLIC Zone, refer to the *SLIC Manager Installation and User Guide*, 310-605807.

The *userlogfile* will define a file with which you can view logged events.

The *QueryChangeInterval* sets the time in seconds that the daemon will poll the Router. The recommended time set here is 10.

Ensure that at the end of every line a semi-colon ';' is used, and that, if several SLIC Zones are created, a comma separates them. Save and exit the file.

9.2.2.3 Installing the SLIC Manager Daemon

With the configuration file edited and a SLIC Zone created, the daemon service can be installed and run. To install the service in Windows NT, open a DOS prompt and go to C:\ibm7190\sdus. Type in `slicd -install`, and the daemon will be installed.

9.2.2.4 Starting the SLIC Manager Daemon

To start the daemon service, select **Start -> Settings -> Control panel** from Windows NT. Double-click on the **Services** icon. Scroll down until you see Vicom SLIC Manager; select and highlight it. You will see two columns to the right to indicate its status. To start the service, select the **Start** button and it will take a few moments to complete. Once it is done, you will see the word `Started` in the Status column. If the Startup column contains the word `Automatic`, no further action is required. If not, select the **Startup** button, and change the Startup Type to `Automatic`. This will have the daemon service start automatically during a reboot. This is shown in Figure 257.

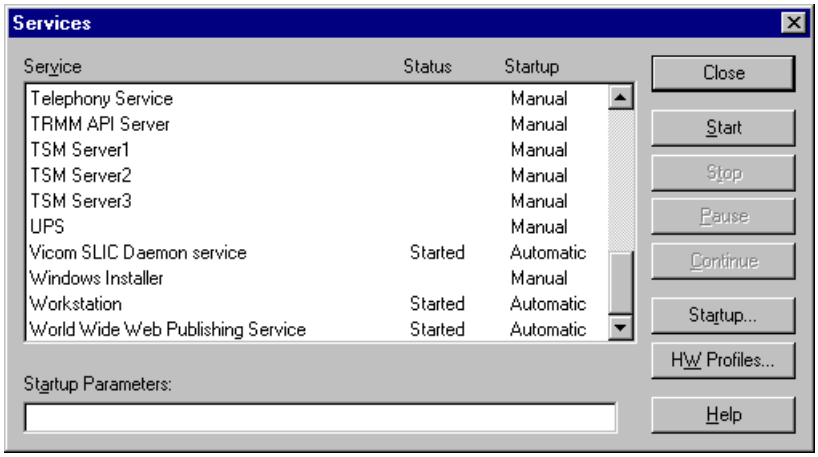

Figure 257. SLIC daemon start up in Windows NT

9.2.3 Starting the SLIC Manager

To start the Vicom SLIC Manager software, select **Start -> Programs -> Vicom -> Vicom SLIC Manager.** The software will load, and a dialog box will appear. In the box with the heading `Hostname`, enter in the name or IP address of the host the daemon service is running. Enter in the SLIC name you entered in when editing the 7190.cfg file. An example is shown in Figure 258.

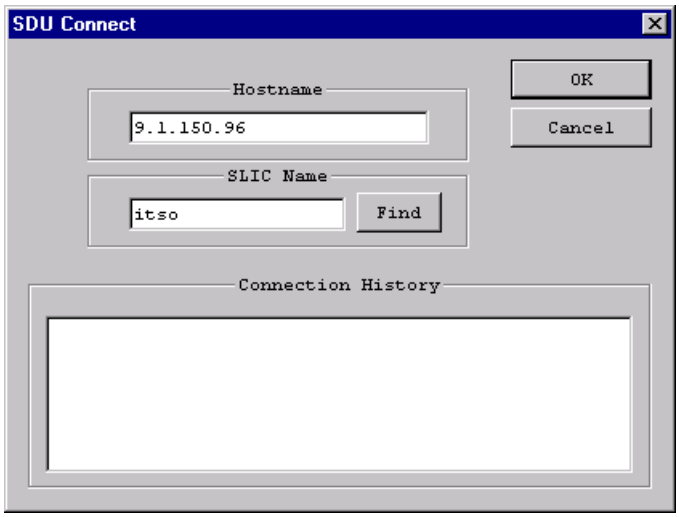

Figure 258. SLIC connection window

Select OK and the software will begin to communicate to the Router. You will notice that the top title bar of your window will now include the host name and SLIC name as in Figure 259.

Figure 259. SLIC Manager title bar

9.3 Using SLIC Manager

You can now look to see that all communications are working properly by going to the toolbar and selecting **Tools -> Control Center**. A dialog box will appear as shown in Figure 260.

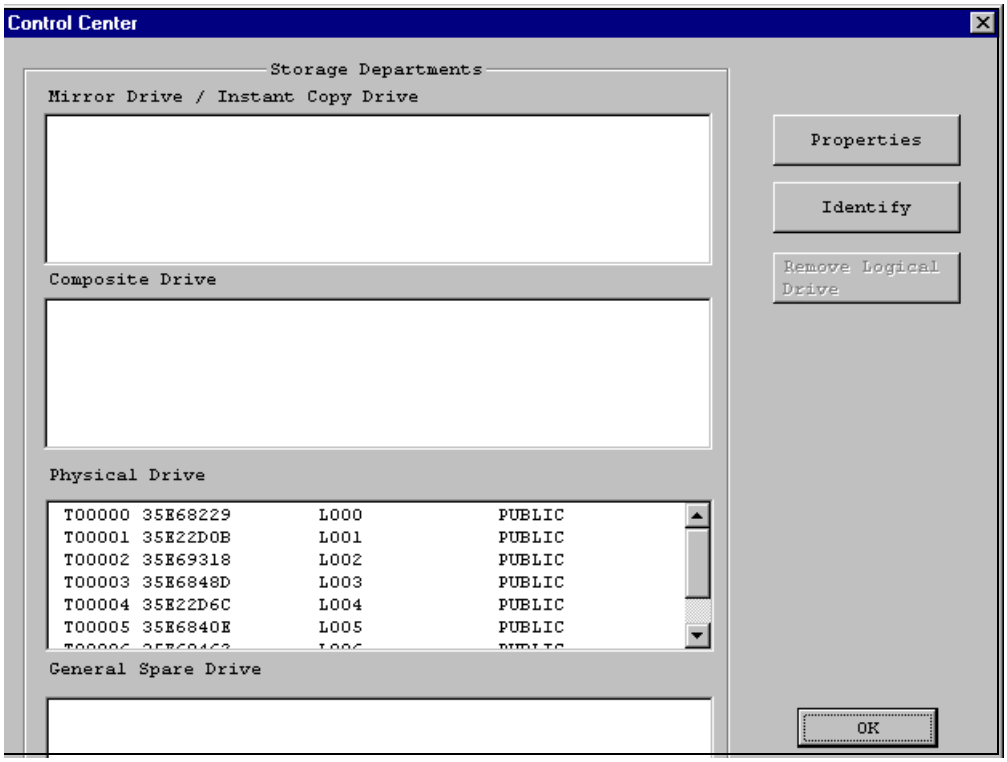

Figure 260. Control Center window

In the Physical Drive box, the drives that are on the SSA loop can be seen. This window will be useful as you start to create mirrors and composite drives, because it provides a summary of all drives.

9.3.1 Drive properties

You can get detailed information on each drive. Select the drive so that it is highlighted and then select **Properties**. A dialog box will appear with the drive's information as in Figure 261.

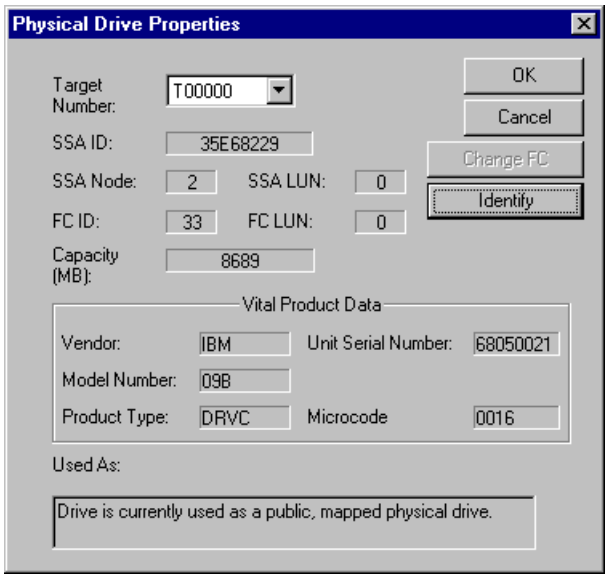

Figure 261. Disk drive properties

Here you can see its SSA attributes, its Fibre Channel attributes and its model type and serial number. By selecting the **Identify** button, the LED on the selected drive will begin to flash.

9.3.2 Router properties

To view the information on the Router, go to the toolbar and select **Properties -> SLIC Properties**. As shown in Figure 262, you will see the serial number of the Router, its ID that was set in SW1, and its supported features.

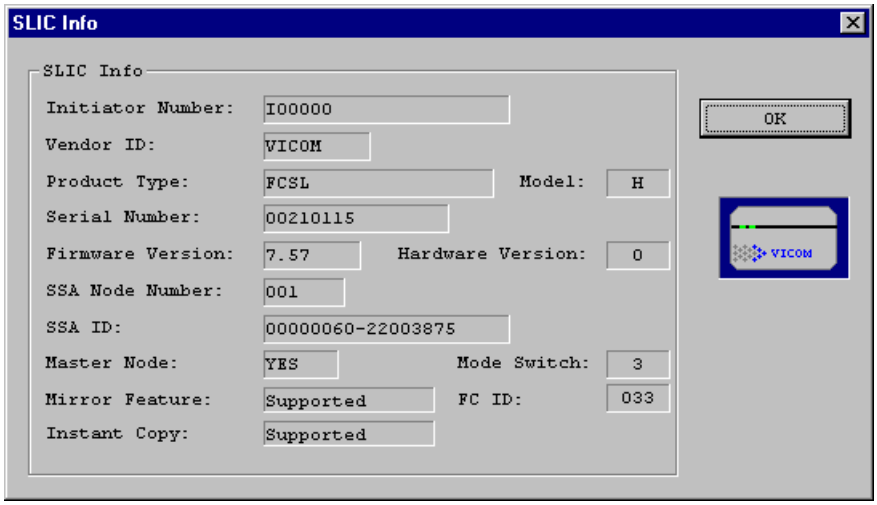

Figure 262. SLIC Router properties

9.3.3 Setting Router to master

As you move through the toolbar, you may notice that most selections have been grayed out. This is due to the fact that the Router is currently in a subordinate role and does not have access to create mirrors or composite drives. This function is done by a Master Router. There can be only one master in a SLIC loop. This is used as more Routers and more disks can be added to the loop. With several Routers in the same loop, there needs to be a requirement where one system acts as the control, and the others will follow and understand any configuration changes that may occur.

To set the Router into a master role select **File -> Program Option** from the top toolbar. You will be presented a dialog box, as shown in Figure 263.

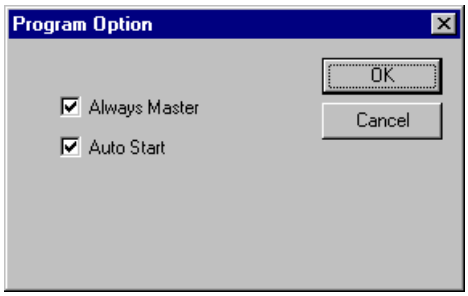

Figure 263. Setting the Router to master

Click in the `Always Master` box so that a check mark appears. Once you select **OK**, the Router will then be set as the master. You will notice that now all options in the toolbar are available and can start to use the features of the SLIC Router.

By placing a check mark in the Auto Start box, the SLIC Manager will automatically connect to the Router defined in the Connection window, as seen in Figure 258 on page 309.

9.3.4 The SignOn drive

When the SLIC zone was created to be used as the communication path, a disk file or partition was created on a specific disk within the SSA loop. As you begin to access the features of the SLIC Router, it should be known which disk was used to create the SLIC zone. This disk is considered to be the SignOn drive.

In the topics 9.4, “Composite drive” on page 314 and 9.5, “Mirror drive” on page 319 we describe creating composite and mirror drives, and you will see that the properties of the individual physical drives may change. As they become part of a logical drive, they take on the properties of this logical drive.

If the SignOn drive is used to create a logical drive, its attributes may change and you may lose the communication path that was created in the SLIC zone. When you select the SignOn drive as a member of a logical drive, a dialog box will be displayed as in Figure 264 to remind you that the attributes of this drive may be affected.

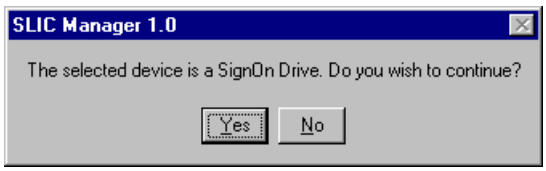

Figure 264. Selecting SignOn drive dialog box

As long as the LUN number of the SignOn drive becomes the LUN of the new logical drive, the communications from the Manager to the Router will not be affected.

Another way to be certain that you do not lose your SignOn drive is not to use the SignOn drive to create logical drives. However, once some logical drives have been created, a SLIC zone can be created to one or more of the newly created logical drives. This logical drive can now be used as the SignOn drive

and the previous drive is now available to be configured without any problems.

9.4 Composite drive

A Composite drive is a large drive that consists of two or more smaller drives. The capacity of the composite drive is an aggregate of the capacities of all the smaller drives that are used to comprise this one large drive.

9.4.1 Creating a composite drive

To create a Composite drive from the SLIC Manager, select **Tools -> Composite Drive Setup Wizard**. A dialog box, Composite Drive List, will appear. Currently, the list will be blank, because there are no Composite drives created. Once there are Composite drives created, you will see a list of the drives. Click on the **Next** button and you will see the Members Selection window, as shown in Figure 265.

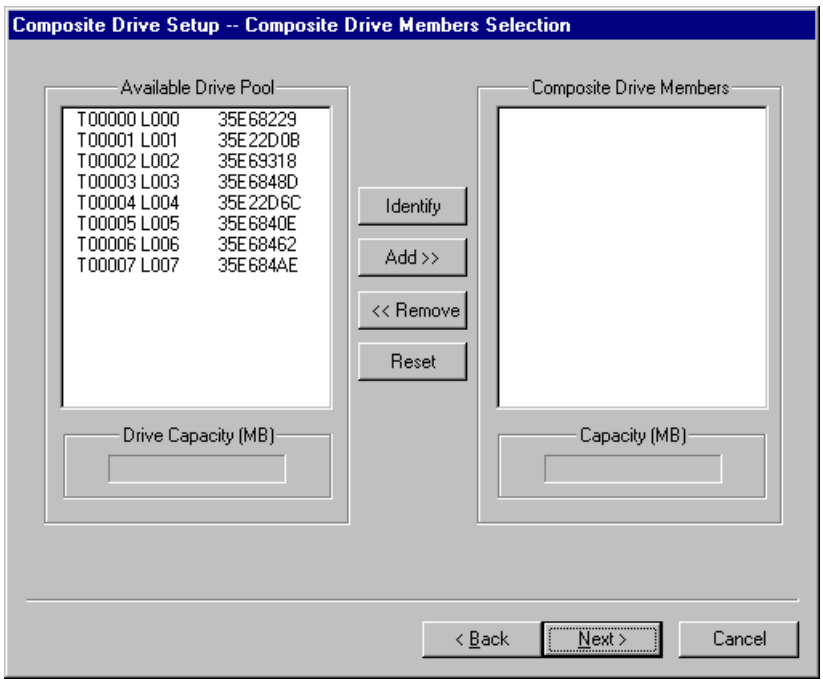

Figure 265. Composite Drive Member Selection window

From the Available Drive Pool list, click on a desired drive and then select the **Add>>** button. The drive name will be added to the Member window. An

asterisk will appear on the left hand side of the drive that was selected in the Available Drive window, to denote that the drive has been selected. Each drive is added one at a time. To remove a drive from the Member window, select the desired drive and click on the **Remove<<** button.

Below each window there is a Drive Capacity box. As a drive is selected, its capacity in megabytes is shown. As you add more member drives to the Member window, the Capacity box will add all drive sizes together to provide a total capacity in megabytes. This is shown in Figure 266.

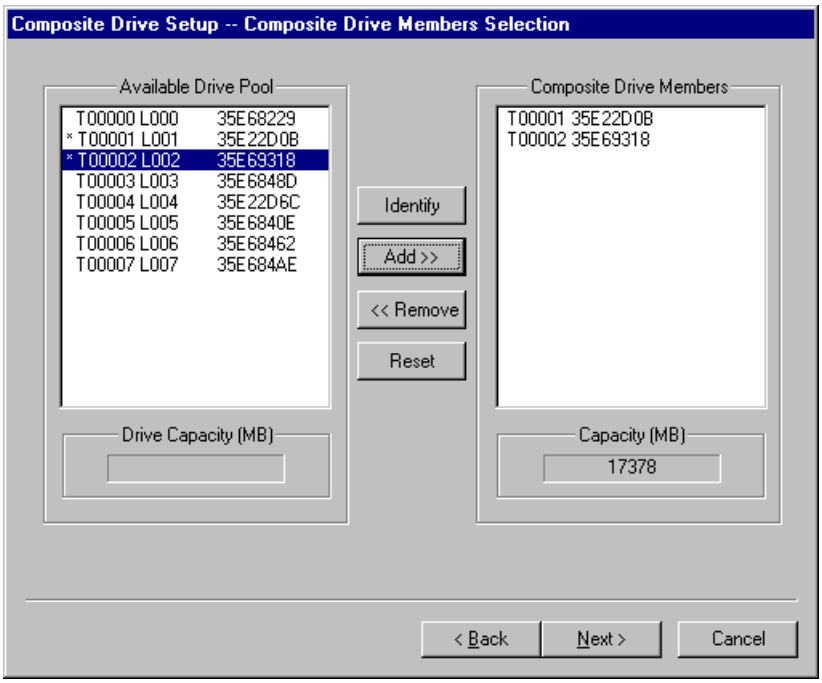

Figure 266. Creating Composite drive from available drives

When all the desired drives are added, select **Next>**. The Assigning Properties window opens, as shown in Figure 267.

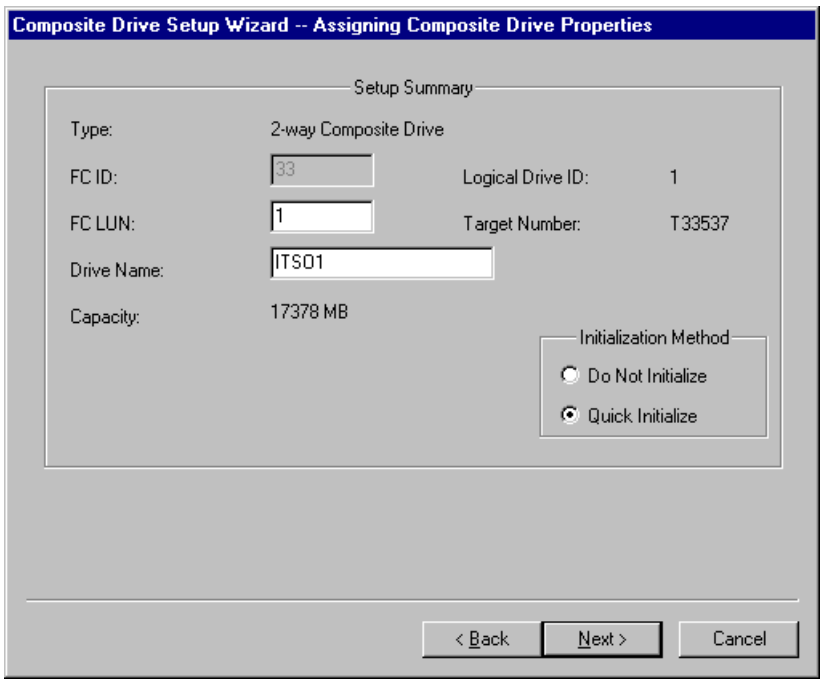

Figure 267. Assigning Composite Drive Properties window

The properties that can be changed are the FC LUN and the Drive Name. There will be a suggested LUN number in this field that can be accepted. If not, simply type in the desired LUN number. The name can also be defined to the Composite drive for easier identification, with a limit of up to eight characters.

The Initialization Method box refers to whether or not to allow the operating system to write its signature on the Composite drive.

Select the **Next>** button and a dialog box will appear, as shown in Figure 268, to allow you to create another Composite drive. Select **Yes** If you would like to create another Composite drive, and the Composite Drive List window opens and the steps described above can be repeated.

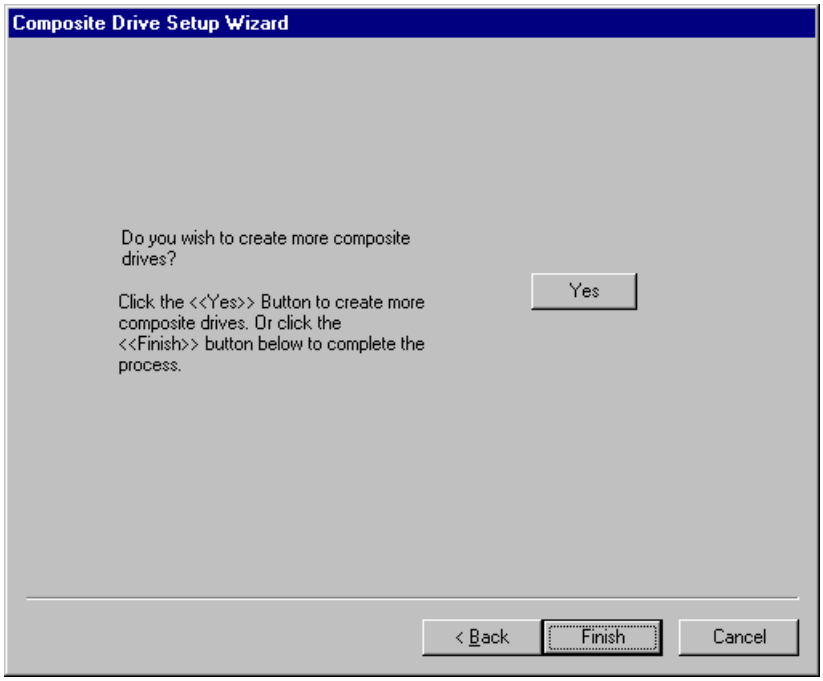

Figure 268. Completing the Composite Drive setup

Select **Finish** when you have created all the desired Composite drives. Up to this point, the configuration has been kept within the SLIC Manager software. When the **Finish** button is selected, the SLIC Manager will now communicate to the Router to complete the process and update the Router to control the drives.

The Host system must rescan for devices, or restart, to be able to see the Composite drive.

9.4.2 Composite drive properties

If you view the Control Center again, by selecting **Tools -> Control Center,** as shown in Figure 269, the newly created Composite drive is listed in the Composite Drive box.

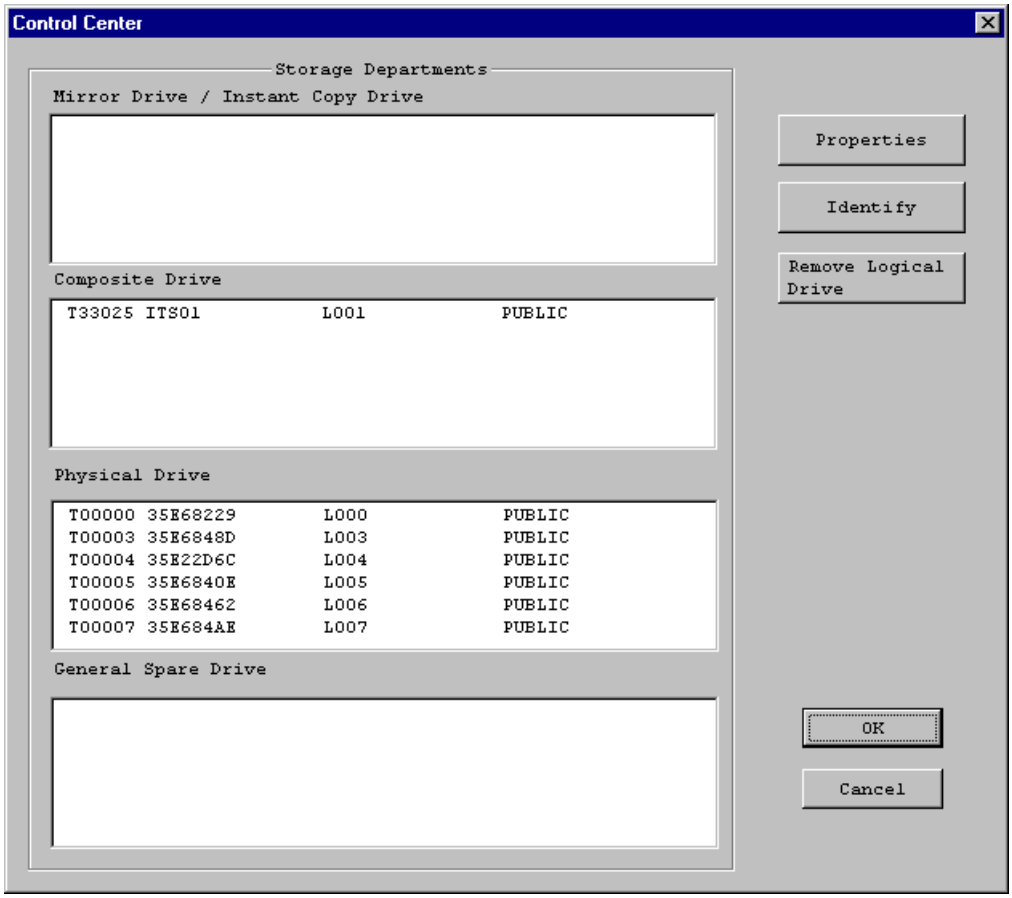

Figure 269. Control Center with Composite drive

Select the Composite drive and then click on the **Properties** button, the Composite Drive Properties dialog box opens, as shown in Figure 270.

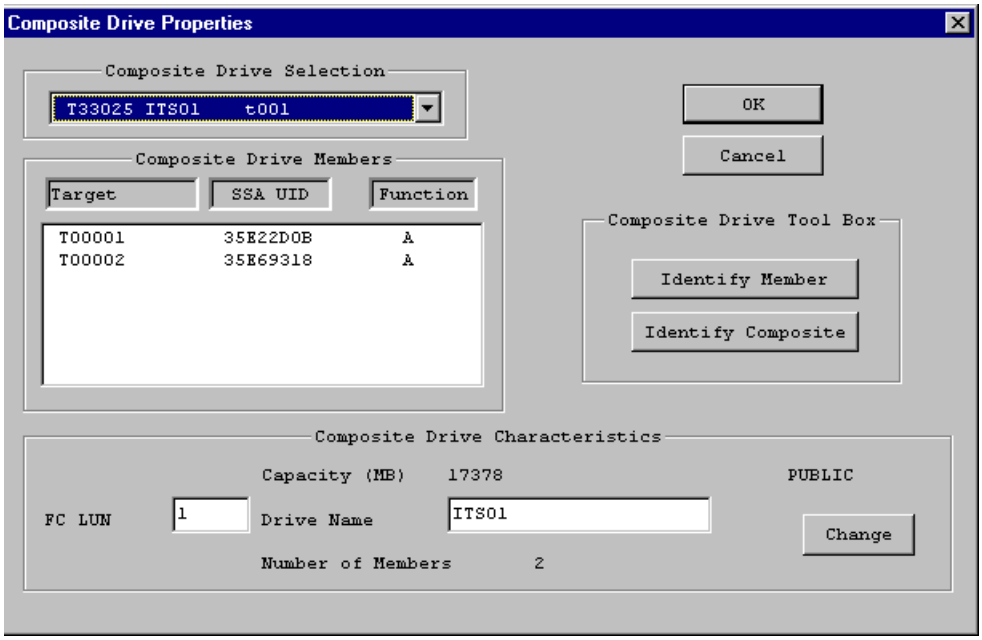

Figure 270. Composite Drive Properties

Here you can find information on the Composite Drive about the member drives that make up the composite, as well as the Composite Drive characteristics. Clicking on the **Identify Composite** button will cause the LED on the actual SSA drives, that belong to the Composite, to flash. If a member drive is selected and the **Identify Member** button is chosen, then the LED only on that drive will flash. In both cases a dialog box will appear to allow you to stop the flashing.

9.5 Mirror drive

A mirror is typically two drives, a 2-way mirror, that contains exactly the same information. The SLIC Router can also support a 1-way or 3-way mirror. A 3-way mirror consists of three drives with the same information. A 1-way mirror is a single drive, or single composite drive, that is used with an Instant Copy Drive that can attach to the single drive mirror to synchronize the data. The Instant Copy Drive can then be split off from the mirror to perform a backup or other action.

The Instant Copy Drive feature can be used with 2-way and 3-way mirrors as well.

9.5.1 Creating a mirror drive

To create a mirror using physical drives, from the toolbar, select **Tools -> Mirror Drive Setup Wizard**. You will see a dialog box, Mirror Drive List, that will be blank. If there were mirror drives created, then it would display the names of the drives. Click on the **Next>>** button and the **Mirror Drive Members Selection** window opens. The window on the left named *Available Drive Pool* contains a list off all drives that are candidates to participate in a mirror drive. Select a drive by highlighting it and click on the **Add>>** button. The drive name will be added to the Member window. An asterisk will appear on the left hand side of the drive just selected in the Available Drive window, to denote that the drive has been selected. A second or third drive can be added to create a 2-way, or 3-way mirror, respectively. Each drive is added one at a time. To remove a drive from the Member window, select the desired drive and click on the **Remove<<** button.

An example of adding two drives to create a 2-way mirror is shown in Figure 271.

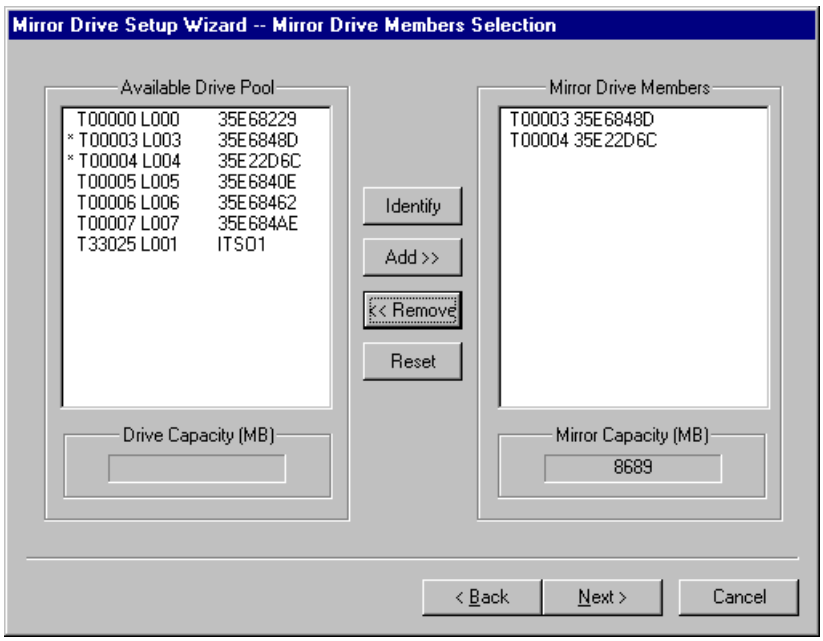

Figure 271. Mirror drive member selection

Below each window there is a Capacity window that will display the size of the available drive, or of the mirror drive. Each drive that participates in a mirror should be of equal capacity. If you select a drive that has a greater capacity

and add it to the mirror, the mirror capacity will still be the smaller of the two, and the rest of the capacity of the larger drive will be unused. For example, if you added a 18 GB drive to the mirror in Figure 271, the Mirror Capacity window would still show the capacity of 8,696 MB. Approximately half of the 18 GB drive will be unused.

After all drives have been added, select **Next>** and you will be able to add a dedicated spare drive to the mirror if desired. Highlight one of the remaining available drives, click on **Add>>** and its name will appear in the Mirror Drive Dedicated Spare window, as shown in Figure 272.

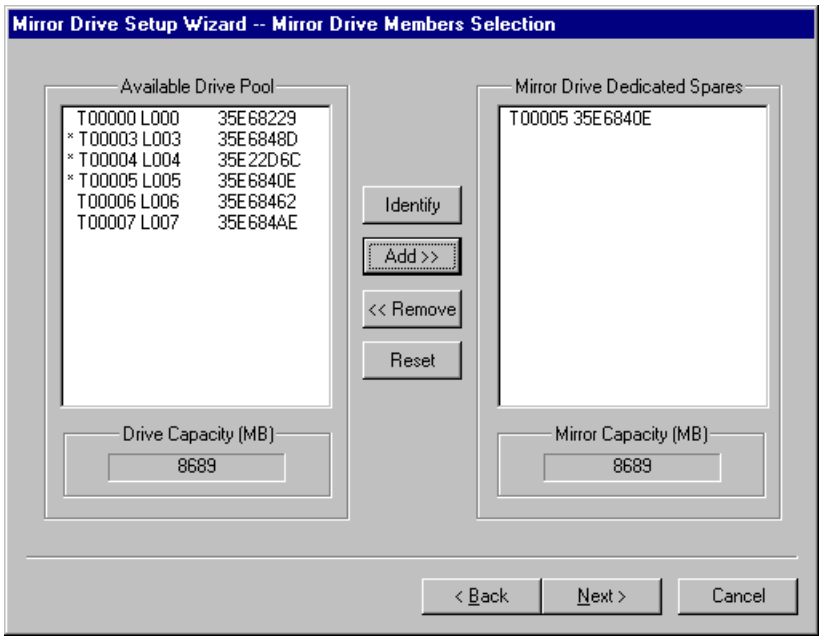

Figure 272. Adding a dedicated spare

Click on the **Next>** button and the properties of the mirror drive can be changed. The properties that can be changed are the FC LUN and the Drive Name. There will be a suggested LUN number in this field that can be accepted. If not, simply type in the desired LUN number. A name can also be defined to the mirror drive for easier identification, with a limit of up to eight characters.

The Initialization Method box refers to whether or not to allow the operating system to write its signature on the Mirror drive.

The Assigning Mirror Drive Properties window is shown in Figure 273.

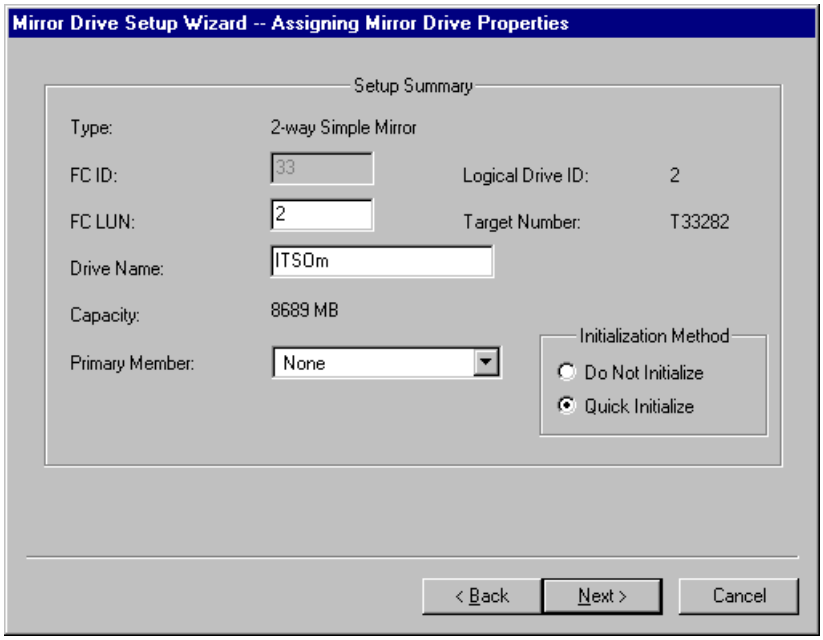

Figure 273. Mirror drive properties

Select the **Next>** button and a dialog box appears to allow you to create another mirror drive. Select **Yes**, if you would like to create another mirror drive, and the Mirror Drive List window opens, and the steps described above can be repeated.

Select **Finish** when you have created all the desired mirror drives. Up to this point, the configuration has been kept within the SLIC Manager software. When the **Finish** button is selected, the SLIC Manager will now communicate to the Router to complete the process and update the Router to control the drives.

If Quick Initialize in the Initialization Method box was selected, the Router will take a short period of time to write the host signature and build the mirror. During this time if you try to communicate to the Router, you may experience a slower than normal response.

9.5.2 Mirror drive properties

If you go to the Control Center window by selecting **Tools -> Control Center**, you will see that the mirror drive is now displayed in the Mirror Drive window. This is shown in Figure 274.

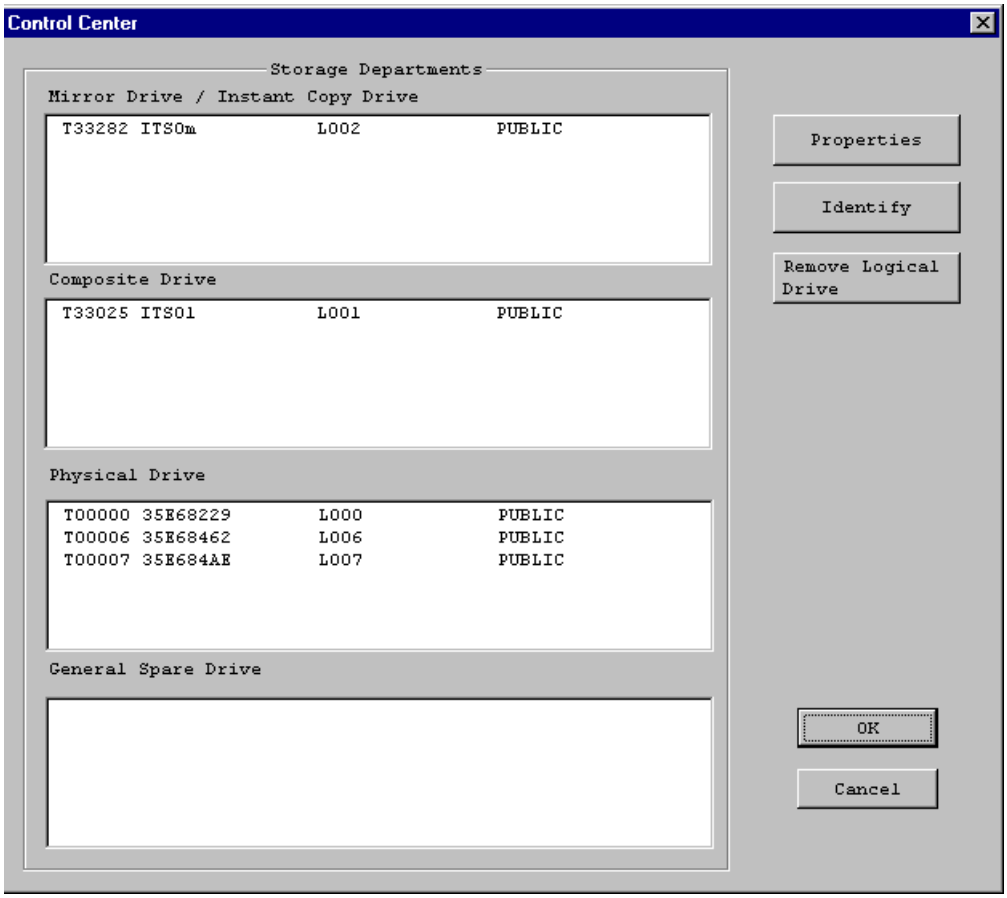

Figure 274. Control Center with Mirror Drive

If you select, and highlight the mirror drive and then click on the **Properties** button, the Mirror Properties window opens, and you can see the information in the mirror drive. Figure 275 shows an example of the properties of the mirror drive.

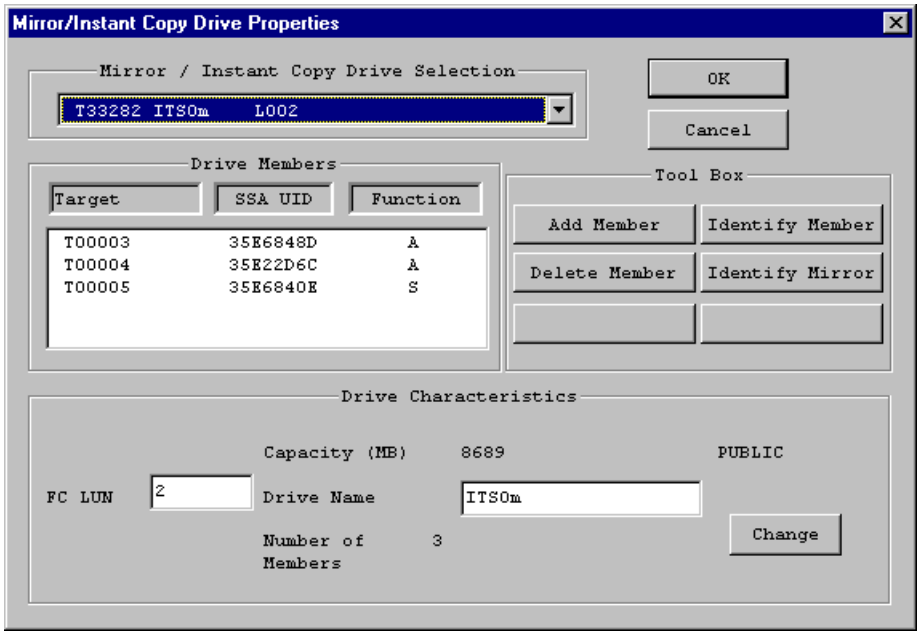

Figure 275. Mirror Drive Properties

Clicking on the **Identify Mirror** button will cause the LED on the actual SSA drives that belong to the mirror, to flash. If a member drive is selected and the **Identify Member** button is chosen, then the LED only on that drive will flash. In both cases a dialog box appears to allow you to stop the flashing.

9.6 Instant Copy drive

Instant Copy is a feature that allows a drive to become part of a mirror, synchronize to the latest data, and then detach from the mirror. The drive can then be used to back up the data or used elsewhere if desired.

9.6.1 Creating an Instant Copy drive

To create an Instant Copy drive, select **Tools -> Instant Copy Drive Setup Wizard**. You will see a dialog box, Instant Copy Drive List, that will be blank. If there were copy drives created, it would display the names of the drives. Click on the **Next>>** button and the **Instant Copy Drive Members Selection** window is displayed. The window on the left named *Available Drive Pool* contains a list off all drives that are candidates to become a copy drive. Select a drive by highlighting it and click on the **Add>>** button. The drive name will be added to the Member window. An asterisk will appear on the left hand side

of the drive just selected in the Available Drive window to denote that the drive has been selected. An example is shown in Figure 276.

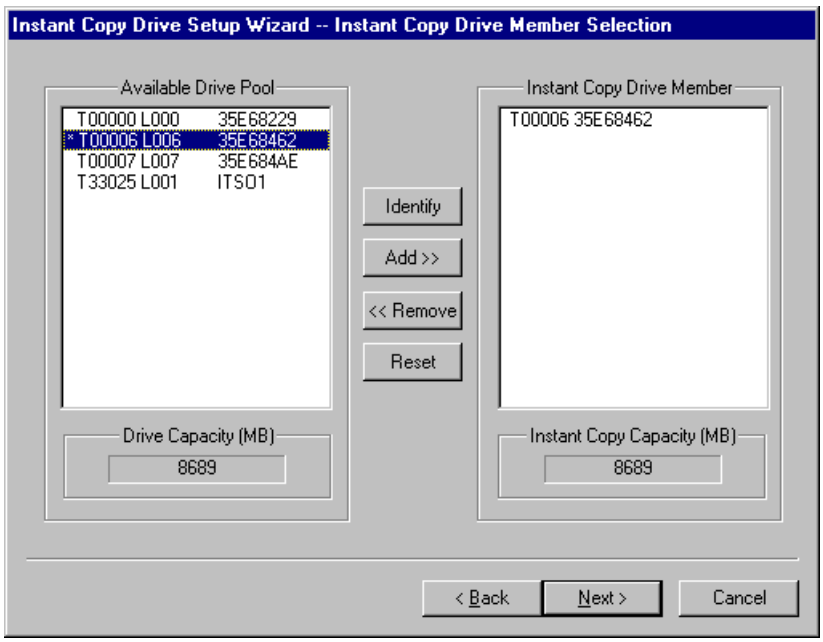

Figure 276. Instant Copy Drive Member Selection

To remove a drive from the Member window, select the desired drive and click on the **Remove<<** button. Below each window there is a Drive Capacity box. As a drive is selected, its capacity in megabytes is shown. Click on the **Next>** button to continue to the Assigning Instant Copy Drive Properties window, as shown in Figure 277.

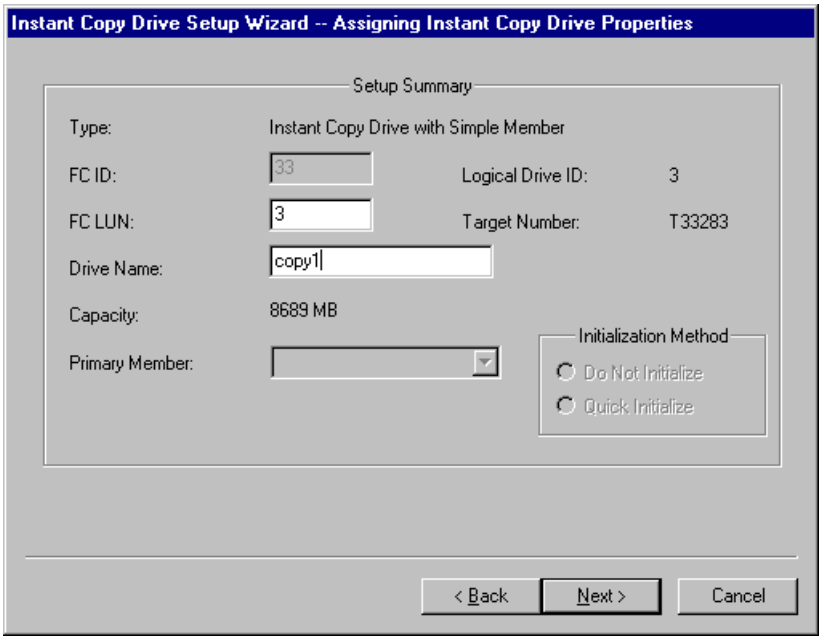

Figure 277. Instant Copy Drive Properties

The properties that can be changed are the FC LUN and the Drive Name. There will be a suggested LUN number in this field that can be accepted. If not, simply type in the desired LUN number. A name can also be defined to the copy drive for easier identification, with a limit of up to eight characters.

Click on the **Next>** button and a dialog box appears to allow you to create another copy drive. Select **Yes** if you would like to create another copy drive and the Instant Copy Drive List window opens, and the steps described above can be repeated.

Select **Finish** when you have created all the desired copy drives. Up to this point, the configuration has been kept within the SLIC Manager software. When the **Finish** button is selected, the SLIC Manager will now communicate to the Router to complete the process and update the Router to control the drives.

9.6.2 Instant copy drive properties

You can go to the Control Center window by selecting **Tools -> Control Center**. The copy drive that was created above can now be seen in the Mirror Drive/Instant Copy window, as shown in Figure 278.

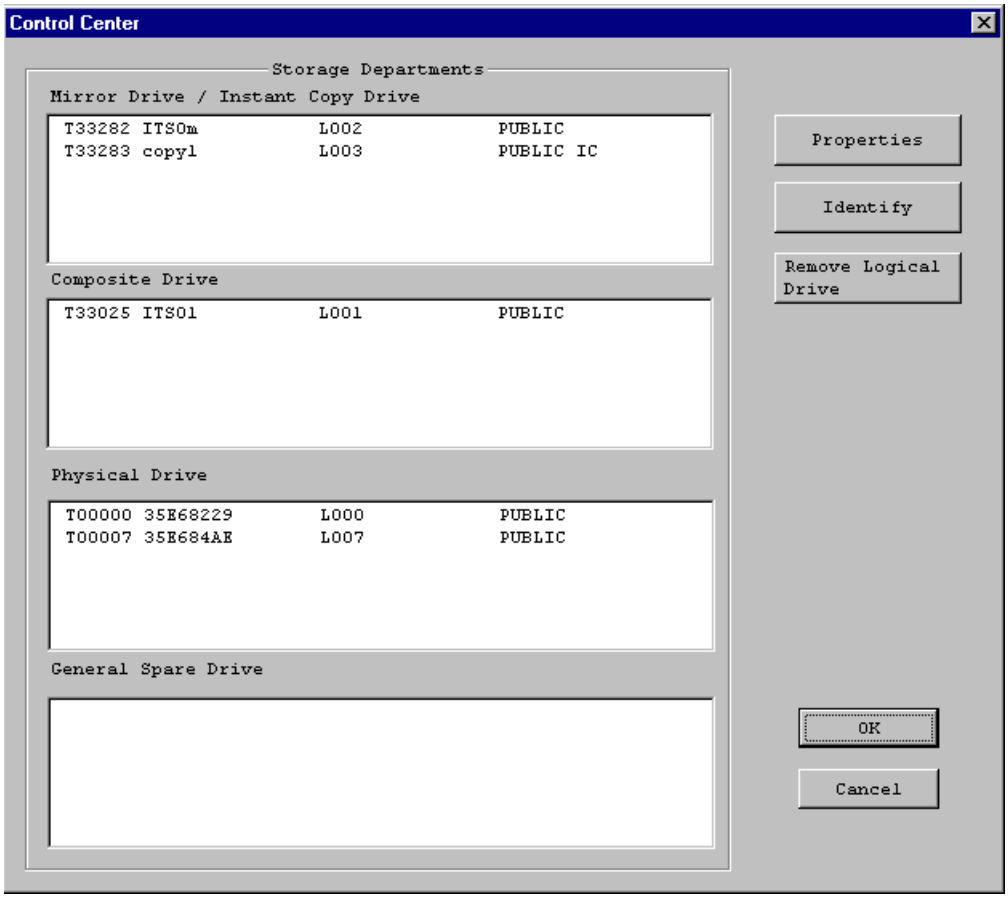

Figure 278. Control Center with Instant Copy Drive

Notice that in the information provided for the copy drive, there is an `IC` included to distinguish between mirror drives and copy drives within this window.

9.6.3 Adding an Instant Copy Drive to a mirror

To add or detach the copy drive from a mirror, you select and highlight the mirror drive, and then click on the **Properties** button. The Mirror Drive Properties window opens, as shown in Figure 275 on page 324. Select the **Add Member** button and the Add Mirror Member window opens, as shown in Figure 279.

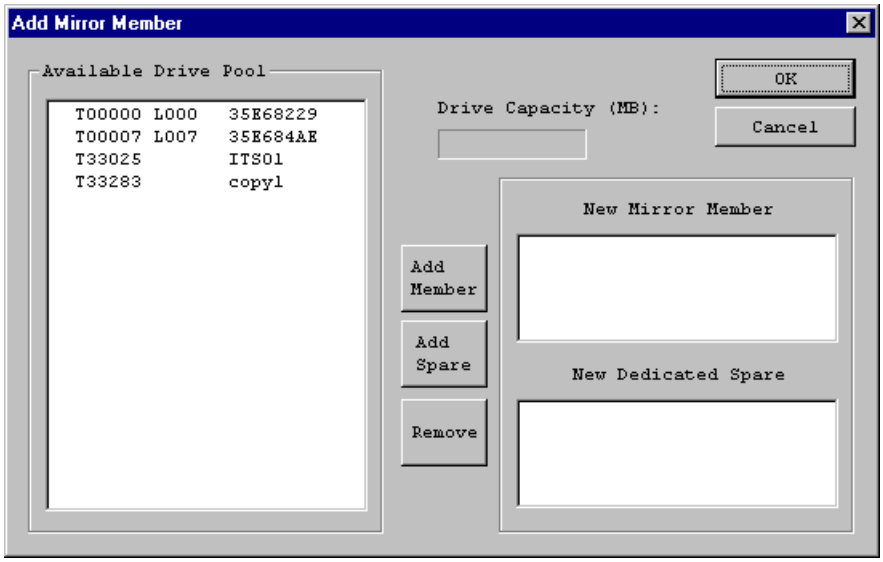

Figure 279. Add Mirror Member display

Select and highlight the copy drive from the Available Drive Pool window, click on the **Add Member** button, and the name of the copy drive will appear in the New Mirror Member window. This is shown in Figure 280.

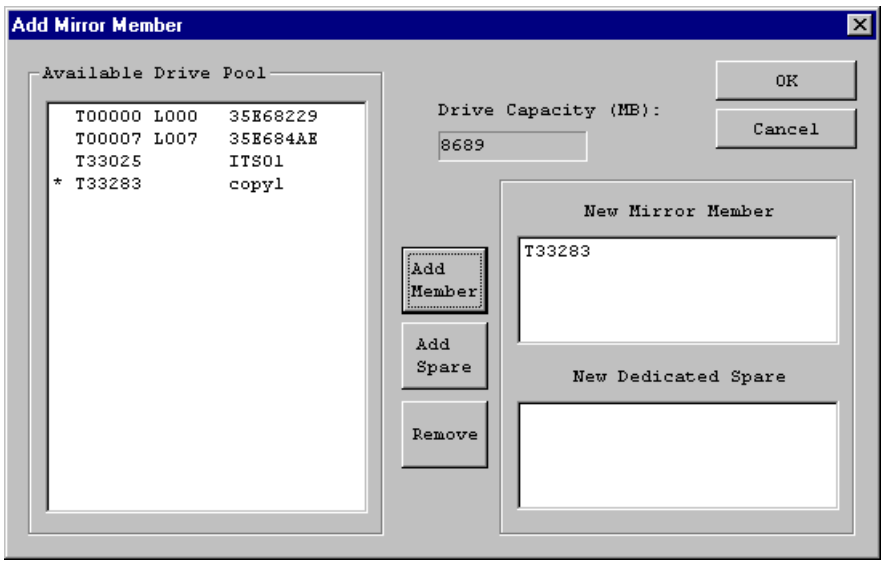

Figure 280. Adding drive members to a mirror

Click on the **OK** button, and the Mirror Drive Properties will now reflect the change, as shown in Figure 281.

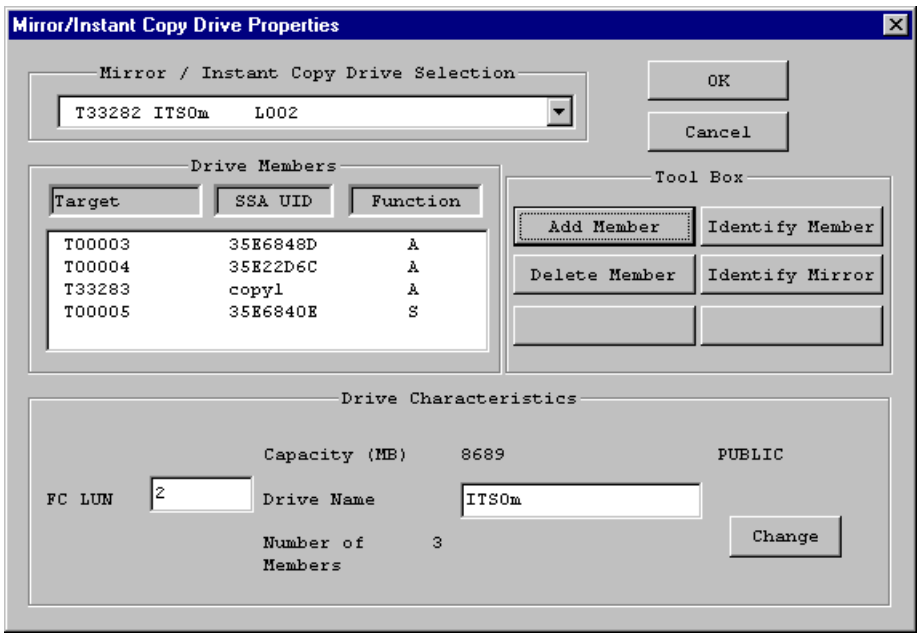

Figure 281. Mirror drive properties with copy drive attached

Click **OK** to complete the process.

9.6.4 Detach Instant Copy Drive from a mirror

To detach, or split off the copy drive from the mirror, the procedure is similar except at the Mirror Drive Properties window, select **Delete Member**. A window will appear that displays all current members of the Mirror. Select the Copy drive, and then delete it from the Mirror. The Copy drive can now be accessed by another host.

9.7 Combining composite and mirroring

The SLIC Manager can also be used to combine the two features of the Router. You can create a mirror drive using composite drives. A mirror can have drive members of different sizes, but the actual mirror capacity will be the smaller of the drive sizes.

9.7.1 Creating a second composite drive

To provide an example of a mirror using only composite drives, another composite drive is required. The example shown in Figure 282 shows that drive 6 and 7 were used to create another composite drive.

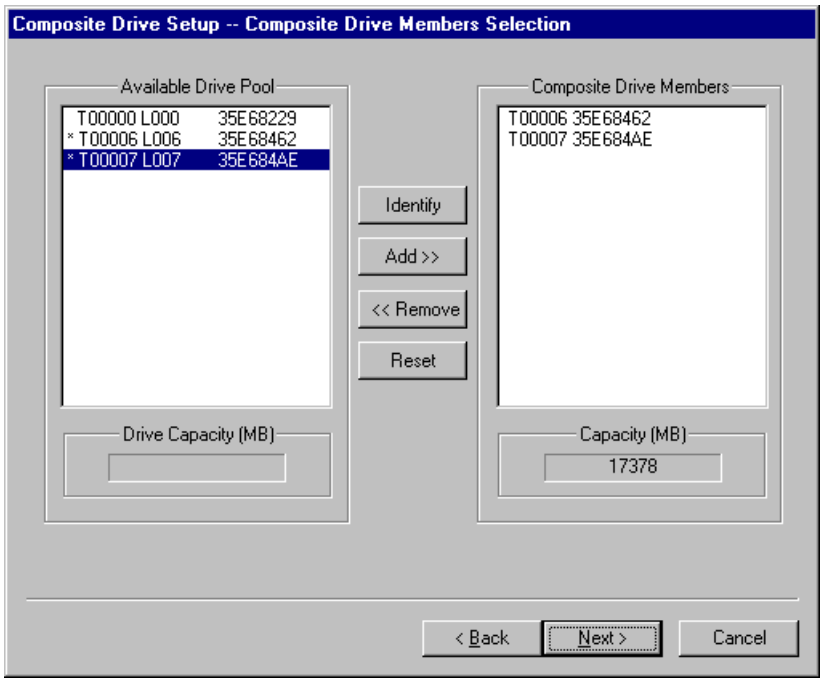

Figure 282. Creating composite drive to be used in a mirror

Follow the steps to create a composite drive as described in 9.4, “Composite drive” on page 314. Once created, you can view the Control Center window by selecting **Tools -> Control Center** from the toolbar. Figure 283 shows that there are now two composite drives.

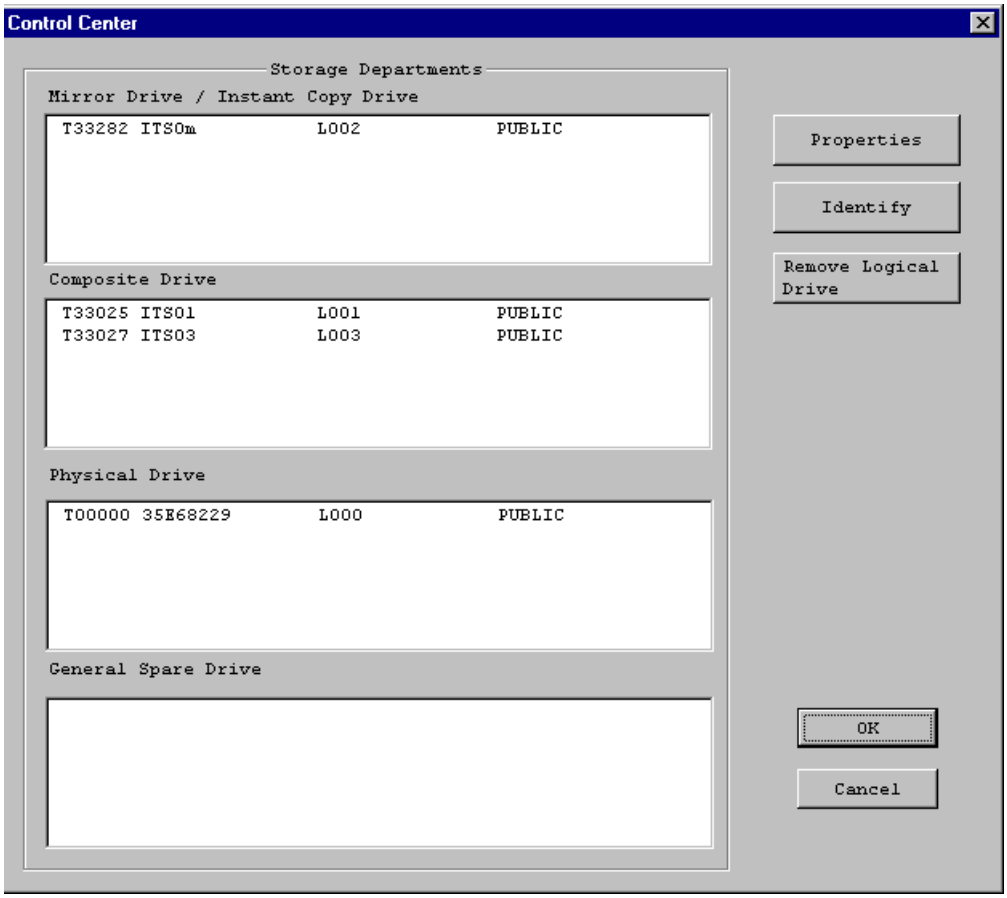

Figure 283. Control Center with two composite drives

9.7.2 Creating the mirror

The mirror can now be created by selecting **Tools -> Mirror Drive Setup Wizard**. When the Member Selection window appears, select the composite drives as members of a mirror. Figure 284 shows where composite drives ‘ITSO1’ and ‘ITSO3’ are selected as members of a mirror.

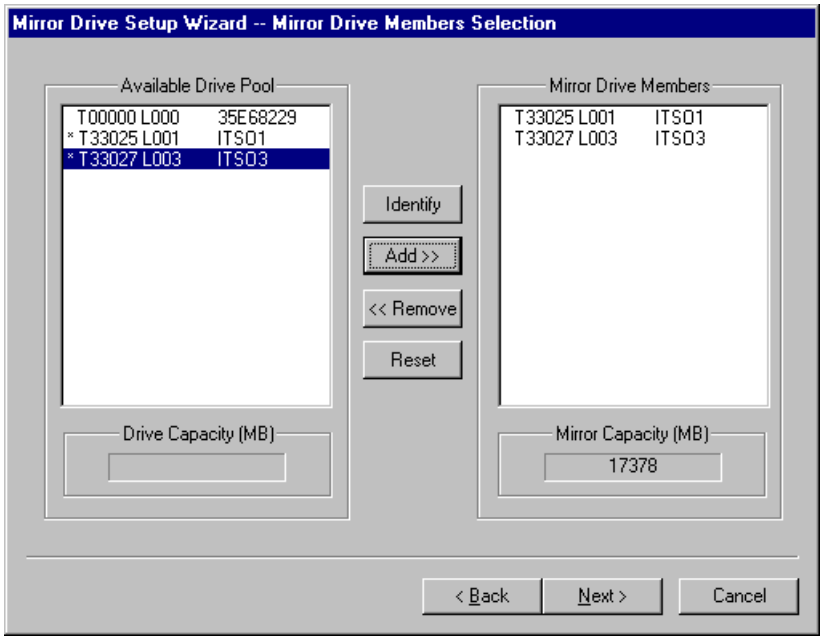

Figure 284. Creating mirror drive from two composite drives

Continue through the Mirror Drive Setup Wizard to complete the process as described in 9.5, “Mirror drive” on page 319.

9.7.3 Viewing mirror drive using composite drives

With the Mirror Drive Setup Wizard completed, you can now view the Control Center window once again, as shown in Figure 285.

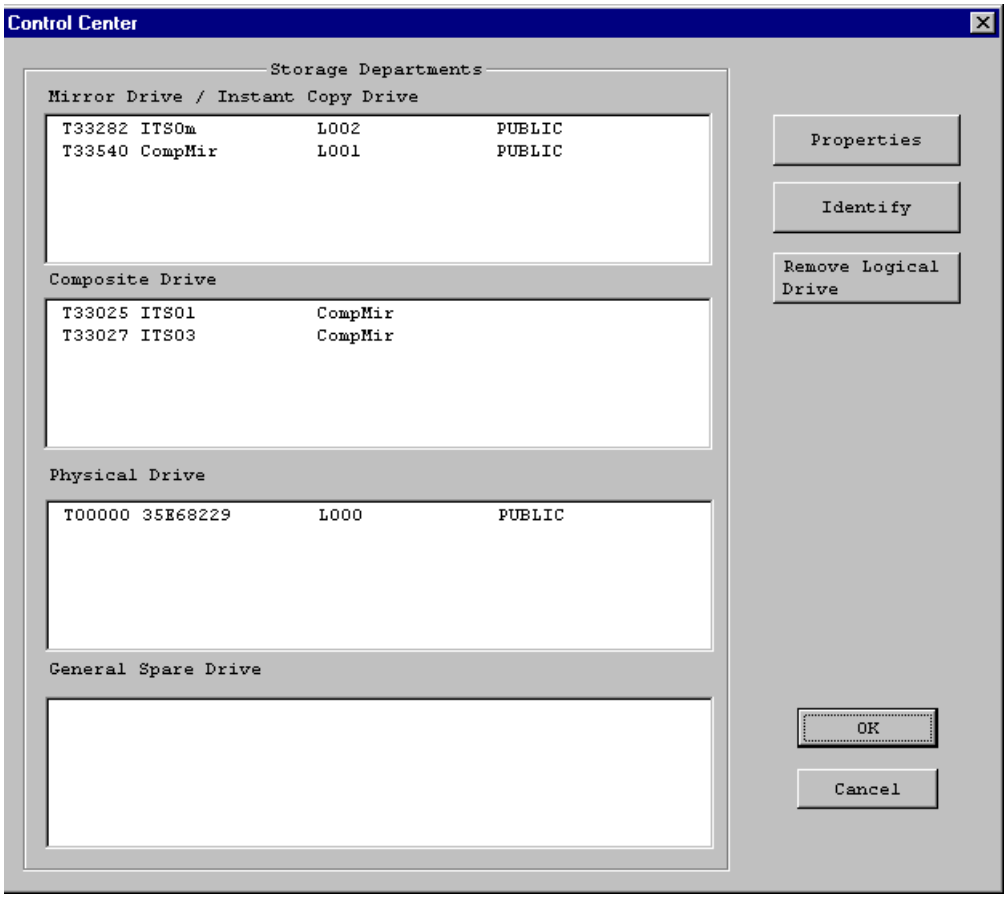

Figure 285. Control Center with mirror drive using two composite drives

In the Mirror Drive window, you can see the new mirror drive that was created above and named *CompMir*. In the Composite Drive window you can see that there are still the two composite drives, but instead of having Fibre Channel LUN numbers assigned to them, they are shown as belonging to a mirror with the name *CompMir*.

You can highlight the CompMir drive and click on the **Properties** button. All the same functions that were described in **9.5.2, "Mirror drive properties" on page 322** are available.

9.8 Reusing logical drives

At some point the composite, mirror, and instant copy logical drives that have created may be no longer required. The logical drive can be removed so that the member drives that made up the logical drive can then be used individually or reconfigured to make new logical drives.

9.8.1 Remove a logical drive

To remove a logical drive, you access the Control Center by selecting **Tools -> Control Center** from the top toolbar. At the Control Center window, select the logical drive (composite, mirror, or copy) that you want to remove. Select the **Remove Logical Drive** button on the right hand side and a dialog box appears that will ask you to confirm that you want to remove the logical drive.

Once it is removed, the member drives will become general spares and will show up in the General Spare Drive window of the Control Center. This is shown in Figure 286.

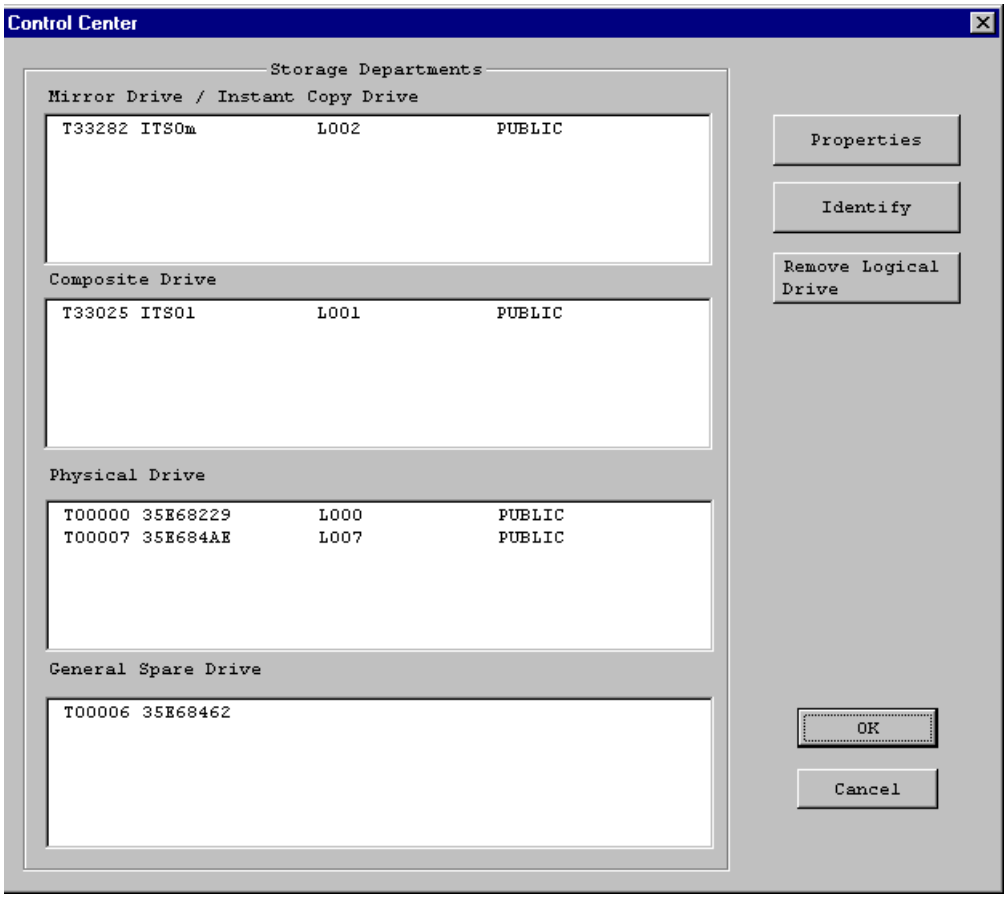

Figure 286. Removing a logical drive

The logical drive that was created as an Instant Copy Drive in 9.6.1, “Creating an Instant Copy drive” on page 324, has been removed and is now a general spare.

9.8.2 Mapping a general spare

You will notice in Figure 286 that the general spare does not have a LUN number assigned to it. To get a new LUN number for this drive, you select the drive and click on the **Properties** button.

The Drive Properties window appears; select the **Change FC** button. A dialog box opens, as shown in Figure 287.

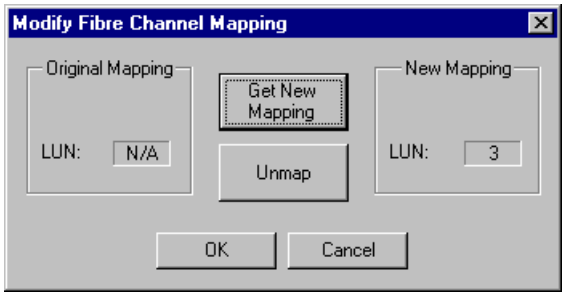

Figure 287. Mapping a general spare

The Original Mapping box will indicate that no LUN was assigned previously. Click on the **Get New Mapping** button, and the next sequential available LUN number will appear in the New Mapping box. Click **OK**. The drive will now appear in the Physical Drive window of the Control Center.

This can also be used to modify the LUN number on an existing mapped drive, as well as remove the LUN number to 'unmap' a drive and create a general spare.

It is not necessary to map a general spare. A general spare can be used to creating a composite, mirror, or copy drive. Mapping a general spare will create a drive that has a LUN number that can then be used by the host.

9.8.3 Removing a mirror containing composite drive

The mirror in this case was made from logical drives on their own. Once the mirror is removed, the composite drives that made up the mirror will return to the Composite Drive window as viewed from the Control Center.

However, since each composite drive had its attributes changed as it became a member of the mirror, it will no longer be mapped. The composite drives will show up as `UnMapped` in the Control Center window. This is shown in Figure 288. The mirror created in 9.7.2, "Creating the mirror" on page 331 was removed.

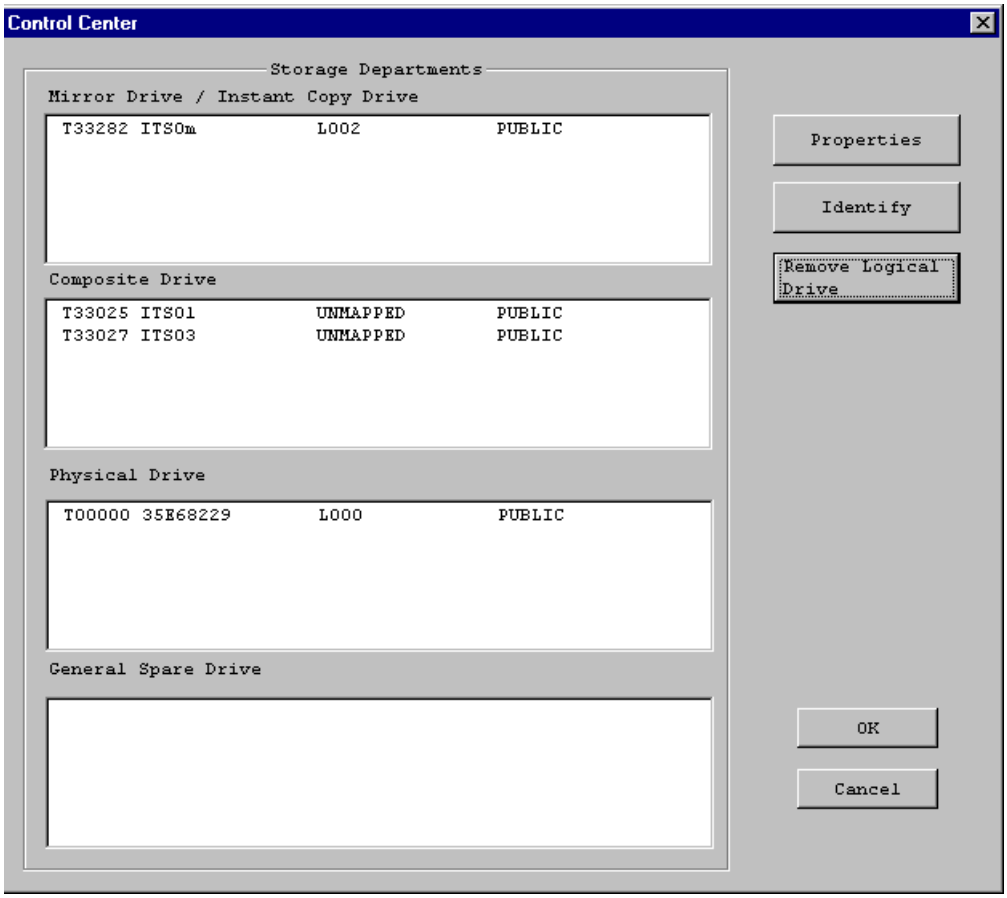

Figure 288. UnMapped composite drives

The existing composite drives 'ITSO1' and 'ITSO3' cannot be mapped or given a new LUN number at this point. Each logical composite drive *must* be removed as well. This will cause the composite drive to be removed and its member drives to become general spares. Once they are general spares, the drives can be used to recreate the composite drive or to create new logical drives.

9.9 Expanding the SLIC system

The SLIC storage system can be expanded to add more SSA disks or more SLIC Routers.

Each SLIC storage system can support up to 64 SSA disks and have 16 Routers.

9.9.1 Adding disk

To add disk to an existing SLIC system is very easy, because they are SSA disks and the rules for SSA disks apply here as well.

If there is a dummy drive in an existing SSA drawer, then it can be replaced by a real drive. If a new drawer has to be added, the SSA cabling is changed to include this new drawer.

Because this is SSA, this can be done *on the fly*. As the SSA loop is broken, the Router will still access all disks due to the structure of the SSA loop. If possible, we recommended that you stop host access and power down the loop. In any case, the rules regarding SSA disks and cabling must be adhered.

As disks are added to an existing loop, the new disks will be recognized. If all disks in the loop are used as single disks (JBOD) and have LUN numbers assigned, the new disks added will have LUN numbers assigned to them automatically. If there are any composite, mirror, instant copy, or spare drives in the loop, then the new disks will not have LUN numbers assigned and become general spares.

9.9.2 Adding Routers

By adding Routers we can increase the amount of storage a host can access and increase throughput. On the rear panel of the Router, there are two Fibre Channel GBIC ports that are available and act as a mini-hub.

You can add a Fibre Channel cable from the second port on the existing Router to one of the ports on the second Router. You are basically daisy-chaining the Routers. But since the ports on the Router act as a hub, an arbitrated loop is created. However, in this scenario there is only one Fibre Channel cable from the host to the Router and it is a single point of failure.

Another option is to add a second Fibre Channel host adapter that will connect to the other Router. This provides a high availability feature, because there are now two paths to the storage system. Software must be used for automatic failover and load balancing between the two Fibre Channel host adapters. Failover also can be done manually if so desired.

On the SSA side, there are a few options available. Each Router can have its own SSA loop so that each one can support 64 SSA disks. In this way,

storage capacity is scalable, because it can be increased by adding more Routers. This is shown in Figure 289.

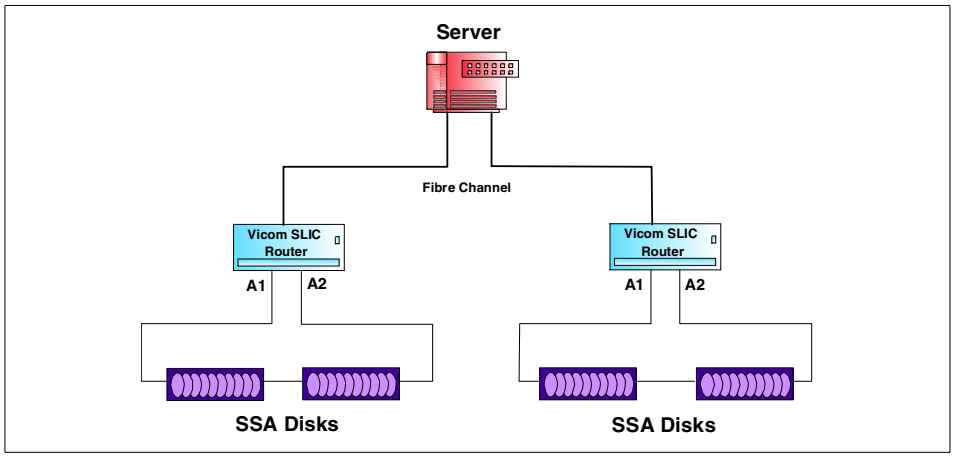

Figure 289. Increasing storage capacity

The other option is to have each additional Router added to the same SSA loop. Throughput to the SSA loop will increase, because each Router can access the disks for multiple simultaneous operations. This configuration is shown in Figure 290.

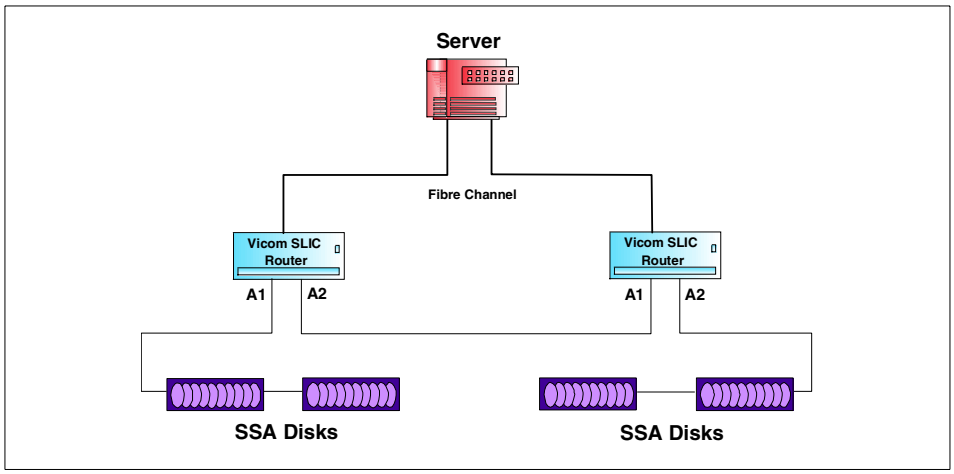

Figure 290. Increasing throughput

9.9.2.1 Installing additional Routers

With the existing Router and storage system powered on:

1. The node map on the new Router must be cleared first.
2. Connect the Router to the existing system with the power off.
3. Set SW2 to mode 3 and set SW1 to an unique Fibre Channel target ID.
4. Power on the new Router
5. When the Status LED on the new Router is on (solid lit), the install is complete.

When the new Router is powered on, communication between the two Routers will occur to query and update the new Router. The new Router will be added as a subordinate so that the first Router will maintain its master status.

Any configuration changes to the storage system is always done on the master. As this is done, the changes are also communicated to the subordinate Routers so that all systems in the loop are aware of what is happening.

9.9.2.2 Using SLIC Manager on additional Routers

Since the master Router does all the work, it is not necessary to use SLIC Manager to view the new Router. However, depending on how the new Router is used this may become a requirement. The same daemon used to connect to the first Router can be used to connect to several Routers.

The configuration file is edited again and the process to name and create a SLIC zone on this new Router can be added within the same file. The SLIC Manager software can now be used to communicate to the new Router.

The SLIC Manager can only communicate to a single Router at a time. Multiple SLIC Manager sessions can be started to communicate to each Router.

9.9.2.3 Master failover

If at some time the Router that is designated as the master within the SLIC storage system fails, the master designation will failover to next nearest Router. This is accomplished within the communications between the Routers and it is done automatically.

When the new Router accepts becoming master, it will maintain the master role if even the failed Router is replaced and rejoins the storage system. The master role can be changed back to the original Router, or to another Router if desired, using the SLIC Manager software.

There is no capability to select a specific ‘failover’ Router.

9.9.3 Adding hosts

The SLIC storage system can be expanded to include more hosts whether they are homogeneous or heterogeneous. It is recommended that as hosts are added, each host is connected to its own and separate SLIC Router.

If more than one host was connected to a single Router, there will be arbitration and performance issues. Also, it would have a single point of failure with the possibility of losing data access to many systems.

9.9.3.1 Homogeneous hosts

If another host of is added and you would like to have both hosts access the same disks, then some sort of access sharing software must be loaded onto both hosts.

If other hosts are added to the storage system and they will not share data, but are connected for storage consolidation, there are a few issues to be considered as the Router does not provide a LUN masking capability.

In UNIX systems, the hosts will see all disk in the loop. But, if the specific volume is not mounted there will be no data integrity problems.

For Windows NT, each host will write its own signature on all available disk. Adding another Windows NT host to the loop will cause problems. To allow a specific Router, and host attached to that Router, access to a specific disk or set of disks, you can set Private Attributes on the disks.

Private Attributes is a setting within SLIC manager that can set a disk to only be accessed by a certain Router and in turn the host attached to that Router.

Note

For more information and operation on the Private Attributes setting, please refer to the SLIC Manager Installation and User Guide, 310-605807

In all cases, if extra control for disk access is required, a third party software, such as Tivoli SANergy, must be used.

9.9.3.2 Heterogeneous hosts

As the Router does not provide for LUN masking, you must use the SLIC Manager Private Attribute setting or a third party software, such as Tivoli

SANergy, to restrict and control host access to the disk. The Private Attributes and Tivoli SANergy can be used together for added control.

Appendix A. Special notices

This publication is intended to help professionals plan and implement a Storage Area Network. The information in this publication is not intended as the specification of any programming interfaces that are provided by the solutions or products mentioned. 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 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.

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

IBM ®	Redbooks
	Redbooks Logo 
AIX	AS/400
AT	DB2
DFSMSrmm	ECKD
Enterprise Storage Server	ESCON
FICON	IBM ,
Magstar	Netfinity
OS/390	OS/400
RS/6000	S/390
Seascape	SP
StorWatch	Ultrastar
Versatile Storage Server	400

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere.,The Power To Manage., Anything. Anywhere.,TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjøbenhavns Sommer - Tivoli A/S.

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, SET Secure Electronic Transaction, 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 B. Related publications

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

B.1 IBM Redbooks

For information on ordering these publications see “How to get IBM Redbooks” on page 351.

- *Designing an IBM Storage Area Network*, SG24-5758
- *Introduction to Storage Area Network, SAN*, SG24-5470
- *IBM Storage Solutions for Server Consolidation*, SG24-5355
- *Implementing the Enterprise Storage Server in Your Environment*, SG24-5420
- *Storage Area Networks: Tape Future In Fabrics*, SG24-5474
- *IBM Enterprise Storage Server*, SG24-5465
- *Introduction to IBM S/390 FICON*, SG24-5176

B.2 IBM Redbooks collections

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

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

B.3 Other resources

These publications are also relevant as further information sources:

- *ESS Web Interface User's Guide for ESS Specialist and ESS Copy Services*, SC26-7346
- *IBM Storage Area Network Data Gateway Installation and User's Guide*, SC26-7304
- *IBM 2109 Model S08 User's Guide*, SC26-7349
- *IBM 2109 Model S08 Switch Service Guide*, SC26-7350
- *IBM 2109 S16 Switch User's Guide*, SC26-7351
- *IBM 2109 S16 Switch Service Guide*, SC26-7352
- *IBM Enterprise Storage Server Configuration Planner*, SC26-7353
- *IBM Enterprise Storage Server Quick Configuration Guide*, SC26-7354
- *IBM SAN Fibre Channel Managed Hub 3534 Service Guide,* SY27-7616
- *IBM SAN Fibre Channel Managed Hub 3534 User's Guide,* GC26-7391
- *IBM Enterprise Storage Server Introduction and Planning Guide, 2105 Models E10, E20, F10 and F20*, GC26-7294
- *IBM Enterprise Storage Server User's Guide, 2105 Models E10, E20, F10 and F20*, SC26-7295
- *IBM Enterprise Storage Server Host Systems Attachment Guide, 2105 Models E10, E20, F10 and F20*, SC26-7296
- *IBM Enterprise Storage Server SCSI Command Reference, 2105 Models E10, E20, F10 and F20*, SC26-7297
- *IBM Enterprise Storage Server System/390 Command Reference, 2105 Models E10, E20, F10 and F20*, SC26-7298
- *IBM Storage Solutions Safety Notices*, GC26-7229
- *Translated External Devices/Safety Information*, SA26-7003
- *Electrical Safety for IBM Customer Engineers*, S229-8124

The following publications can be ordered at `http://www.vicom.com`

- *SLIC Router FC-SL Installation and User Guide*, 310-605759
- *SLIC Manager Installation and User Guide*, 310-605807

B.4 Referenced Web sites

These Web sites are also relevant as further information sources:

- http://www.storage.ibm.com/ibmsan/index.htm
  IBM Enterprise SAN
- http://www.storage.ibm.com/hardsoft/products/fchub/fchub.htm
  IBM Fibre Channel Storage HUB
- http://www.pc.ibm.com/ww/netfinity/san
  IBM Storage Area Networks: Nefinity Servers
- http://www.storage.ibm.com/hardsoft/products/fcswitch/fcswitch.htm
  IBM SAN Fibre Channel Switch
- http://www.storage.ibm.com/hardsoft/products/fchub/fcmhub.htm
  IBM SAN Fibre Channel Managed Hub
- http://www.storage.ibm.com/hardsoft/products/sangateway/supserver.htm
  IBM SAN Data Gateway
- http://www.storage.ibm.com/hardsoft/products/tape/ro3superserver.htm
  IBM SAN Data Gateway Router
- http://www.storage.ibm.com/hardsoft/products/fcss/fcss.htm
  IBM Fibre Channel RAID Storage Server
- http://www.storage.ibm.com/hardsoft/products/ess/ess.htm
  Enterprise Storage Server
- http://www.brocade.com
  Brocade Communications Systems, Inc.
- http://www.fibrechannel.com
  Fibre Channel Industry Association
- http://www.mcdata.com
  McDATA Corporation
- http://www.pathlight.com
  Pathlight
- http://www.sanergy.com
  Tivoli SANergy
- http://www.snia.org
  Storage Networking Industry Association
- http://www.tivoli.com
  Tivoli Systems
- http://www.t11.org
  Technical Committee T11
- http://www.vicom.com
  Vicom Systems
- http://www.vixel.com
  Vixel

How to get IBM Redbooks

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

- **Redbooks Web Site** ibm.com/redbooks

  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 IBM Redbooks fax order form to:

	e-mail address
In United States or Canada	pubscan@us.ibm.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 Redbooks 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.

Glossary

8B/10B A data encoding scheme developed by IBM, translating byte-wide data to an encoded 10-bit format. Fibre Channel's FC-1 level defines this as the method to be used to encode and decode data transmissions over the Fibre channel.

Adapter A hardware unit that aggregates other I/O units, devices or communications links to a system bus.

ADSM Adstar Distributed Storage Manager

Agent (1) In the client-server model, the part of the system that performs information preparation and exchange on behalf of a client or server application. (2) In SNMP, the word agent refers to the managed system. See also: Management Agent

AIT Advanced Intelligent Tape - A magnetic tape format by Sony that uses 8mm cassettes, but is only used in specific drives.

AL See Arbitrated Loop

ANSI American National Standards Institute - The primary organization for fostering the development of technology standards in the United States. The ANSI family of Fibre Channel documents provide the standards basis for the Fibre Channel architecture and technology. See FC-PH

Arbitration The process of selecting one respondent from a collection of several candidates that request service concurrently.

Arbitrated Loop A Fibre Channel interconnection technology that allows up to 126 participating node ports and one participating fabric port to communicate.

ATL Automated Tape Library - Large scale tape storage system, which uses multiple tape drives and mechanisms to address 50 or more cassettes.

ATM Asynchronous Transfer Mode - A type of packet switching that transmits fixed-length units of data.

Backup A copy of computer data that is used to recreate data that has been lost, mislaid, corrupted, or erased. The act of creating a copy of computer data that can be used to recreate data that has been lost, mislaid, corrupted or erased.

Bandwidth Measure of the information capacity of a transmission channel.

Bridge (1) A component used to attach more than one I/O unit to a port. (2) A data communications device that connects two or more networks and forwards packets between them. The bridge may use similar or dissimilar media and signaling systems. It operates at the data link level of the OSI model. Bridges read and filter data packets and frames.

Bridge/Router A device that can provide the functions of a bridge, router or both concurrently. A bridge/router can route one or more protocols, such as TCP/IP, and bridge all other traffic. See also: Bridge, Router

Broadcast Sending a transmission to all N_Ports on a fabric.

Channel A point-to-point link, the main task of which is to transport data from one point to another.

Channel I/O A form of I/O where request and response correlation is maintained through some form of source, destination and request identification.

CIFS Common Internet File System

Class of Service A Fibre Channel frame delivery scheme exhibiting a specified set of delivery characteristics and attributes.

Class-1 A class of service providing dedicated connection between two ports with confirmed delivery or notification of non-deliverability.

Class-2 A class of service providing a frame switching service between two ports with confirmed delivery or notification of non-deliverability.

Class-3 A class of service providing frame switching datagram service between two ports or a multicast service between a multicast originator and one or more multicast recipients.

Class-4 A class of service providing a fractional bandwidth virtual circuit between two ports with confirmed delivery or notification of non-deliverability.

Class-6 A class of service providing a multicast connection between a multicast originator and one or more multicast recipients with confirmed delivery or notification of non-deliverability.

Client A software program used to contact and obtain data from a *server* software program on another computer -- often across a great distance. Each *client* program is designed to work specifically with one or more kinds of server programs and each server requires a specific kind of client program.

Client/Server The relationship between machines in a communications network. The client is the requesting machine, the server the supplying machine. Also used to describe the information management relationship between software components in a processing system.

Cluster A type of parallel or distributed system that consists of a collection of interconnected whole computers and is used as a single, unified computing resource.

Coaxial Cable A transmission media (cable) used for high speed transmission. It is called *coaxial* because it includes one physical channel that carries the signal surrounded (after a layer of insulation) by another concentric physical channel, both of which run along the same axis. The inner channel carries the signal and the outer channel serves as a ground.

Controller A component that attaches to the system topology through a channel semantic protocol that includes some form of request/response identification.

CRC Cyclic Redundancy Check - An error-correcting code used in Fibre Channel.

DASD Direct Access Storage Device - any on-line storage device: a disc, drive or CD-ROM.

DAT Digital Audio Tape - A tape media technology designed for very high quality audio recording and data backup. DAT cartridges look like audio cassettes and are often used in mechanical auto-loaders. typically, a DAT cartridge provides 2GB of storage. But new DAT systems have much larger capacities.

Data Sharing A SAN solution in which files on a storage device are shared between multiple hosts.

Datagram Refers to the Class 3 Fibre Channel Service that allows data to be sent rapidly to multiple devices attached to the fabric, with no confirmation of delivery.

dB Decibel - a ratio measurement distinguishing the percentage of signal attenuation between the input and output power. Attenuation (loss) is expressed as dB/km

Disk Mirroring A fault-tolerant technique that writes data simultaneously to two hard disks using the same hard disk controller.

Disk Pooling A SAN solution in which disk storage resources are pooled across multiple hosts rather than be dedicated to a specific host.

DLT Digital Linear Tape - A magnetic tape technology originally developed by Digital Equipment Corporation (DEC) and now sold by Quantum. DLT cartridges provide storage capacities from 10 to 35GB.

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

ECL Emitter Coupled Logic - The type of transmitter used to drive copper media such as Twinax, Shielded Twisted Pair, or Coax.

Enterprise Network A geographically dispersed network under the auspices of one organization.

Entity In general, a real or existing thing from the Latin ens, or being, which makes the distinction between a thing's existence and it qualities. In programming, engineering and probably many other contexts, the word is used to identify units, whether concrete things or abstract ideas, that have no ready name or label.

ESCON Enterprise System Connection

Exchange A group of sequences which share a unique identifier. All sequences within a given exchange use the same protocol. Frames from multiple sequences can be multiplexed to prevent a single exchange from consuming all the bandwidth. See also: Sequence

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.

Fabric Fibre Channel employs a fabric to connect devices. A fabric can be as simple as a single cable connecting two devices. The term is most often used to describe a more complex network utilizing hubs, switches and gateways.

Fabric Login Fabric Login (FLOGI) is used by an N_Port to determine if a fabric is present and, if so, to initiate a session with the fabric by exchanging service parameters with the fabric. Fabric Login is performed by an N_Port following link initialization and before communication with other N_Ports is attempted.

FC Fibre Channel

FC-0 Lowest level of the Fibre Channel Physical standard, covering the physical characteristics of the interface and media

FC-1 Middle level of the Fibre Channel Physical standard, defining the 8B/10B encoding/decoding and transmission protocol.

FC-2 Highest level of the Fibre Channel Physical standard, defining the rules for signaling protocol and describing transfer of frame, sequence and exchanges.

FC-3 The hierarchical level in the Fibre Channel standard that provides common services such as striping definition.

FC-4 The hierarchical level in the Fibre Channel standard that specifies the mapping of upper-layer protocols to levels below.

FCA Fiber Channel Association.

FC-AL Fibre Channel Arbitrated Loop - A reference to the Fibre Channel Arbitrated Loop standard, a shared gigabit media for up to 127 nodes, one of which may be attached to a switch fabric. See also: Arbitrated Loop.

FC-CT Fibre Channel common transport protocol

FC-FG Fibre Channel Fabric Generic - A reference to the document (ANSI X3.289-1996) which defines the concepts, behavior and characteristics of the Fibre Channel Fabric along with suggested partitioning of the 24-bit address space to facilitate the routing of frames.

FC-FP Fibre Channel HIPPI Framing Protocol - A reference to the document (ANSI X3.254-1994) defining how the HIPPI framing protocol is transported via the fibre channel

FC-GS Fibre Channel Generic Services -A reference to the document (ANSI X3.289-1996) describing a common transport protocol used to communicate with the server functions, a full X500 based directory service, mapping of the Simple Network Management Protocol (SNMP) directly to the Fibre Channel, a time server and an alias server.

FC-LE Fibre Channel Link Encapsulation - A reference to the document (ANSI X3.287-1996) which defines how IEEE 802.2 Logical Link Control (LLC) information is transported via the Fibre Channel.

FC-PH A reference to the Fibre Channel Physical and Signaling standard ANSI X3.230, containing the definition of the three lower levels (FC-0, FC-1, and FC-2) of the Fibre Channel.

FC-PLDA Fibre Channel Private Loop Direct Attach - See PLDA.

FC-SB Fibre Channel Single Byte Command Code Set - A reference to the document (ANSI X.271-1996) which defines how the ESCON command set protocol is transported using the fibre channel.

FC-SW Fibre Channel Switch Fabric - A reference to the ANSI standard under development that further defines the fabric behavior described in FC-FG and defines the communications between different fabric elements required for those elements to coordinate their operations and management address assignment.

FC Storage Director See SAN Storage Director

FCA Fibre Channel Association - a Fibre Channel industry association that works to promote awareness and understanding of the Fibre Channel technology and its application and provides a means for implementers to support the standards committee activities.

FCLC Fibre Channel Loop Association - an independent working group of the Fibre Channel Association focused on the marketing aspects of the Fibre Channel Loop technology.

FCP Fibre Channel Protocol - the mapping of SCSI-3 operations to Fibre Channel.

Fiber Optic Refers to the medium and the technology associated with the transmission of information along a glass or plastic wire or fiber.

Fibre Channel A technology for transmitting data between computer devices at a data rate of up to 4 Gb/s. It is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

FICON Fibre Connection - A next-generation I/O solution for IBM S/390 parallel enterprise server.

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

FLOGI See Fabric Log In

Frame A linear set of transmitted bits that define the basic transport unit. The frame is the most basic element of a message in Fibre Channel communications, consisting of a 24-byte header and zero to 2112 bytes of data. See also: Sequence

FSP Fibre Channel Service Protocol - The common FC-4 level protocol for all services, transparent to the fabric type or topology.

Full-Duplex A mode of communications allowing simultaneous transmission and reception of frames.

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

Gateway A node on a network that interconnects two otherwise incompatible networks.

GBIC GigaBit Interface Converter - Industry standard transceivers for connection of Fibre Channel nodes to arbitrated loop hubs and fabric switches.

Gigabit One billion bits, or one thousand megabits.

GLM Gigabit Link Module - a generic Fibre Channel transceiver unit that integrates the key functions necessary for installation of a Fibre channel media interface on most systems.

Half-Duplex A mode of communications allowing either transmission or reception of frames at any point in time, but not both (other than link control frames which are always permitted).

Hardware The mechanical, magnetic and electronic components of a system, e.g., computers, telephone switches, terminals and the like.

HBA Host Bus Adapter

HIPPI High Performance Parallel Interface - An ANSI standard defining a channel that transfers data between CPUs and from a CPU to disk arrays and other peripherals.

HMMP HyperMedia Management Protocol

HMMS HyperMedia Management Schema - the definition of an implementation-independent, extensible, common data description/schema allowing data from a variety of sources to be described and accessed in real time regardless of the source of the data. See also: WEBM, HMMP

HSM Hierarchical Storage Management - A software and hardware system that moves files from disk to slower, less expensive storage media based on rules and observation of file activity. Modern HSM systems move files from magnetic disk to optical disk to magnetic tape.

HUB A Fibre Channel device that connects nodes into a logical loop by using a physical star topology. Hubs will automatically recognize an active node and insert the node into the loop. A node that fails or is powered off is automatically removed from the loop.

HUB Topology see Loop Topology

Hunt Group A set of associated Node Ports (N_Ports) attached to a single node, assigned a special identifier that allows any frames containing this identifier to be routed to any available Node Port (N_Port) in the set.

In-Band Signaling Signaling that is carried in the same channel as the information.

Information Unit A unit of information defined by an FC-4 mapping. Information Units are transferred as a Fibre Channel Sequence.

Intermix A mode of service defined by Fibre Channel that reserves the full Fibre Channel bandwidth for a dedicated Class 1 connection, but also allows connection-less Class 2 traffic to share the link if the bandwidth is available.

I/O Input/output

IP Internet Protocol

IPI Intelligent Peripheral Interface

Isochronous Transmission Data transmission which supports network-wide timing requirements. A typical application for isochronous transmission is a broadcast environment which needs information to be delivered at a predictable time.

JBOD Just a bunch of disks.

Jukebox A device that holds multiple optical disks and one or more disk drives, and can swap disks in and out of the drive as needed.

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.

LAN See Local Area Network - A network covering a relatively small geographic area (usually not larger than a floor or small building). Transmissions within a Local Area Network are mostly digital, carrying data among stations at rates usually above one megabit/s.

Latency A measurement of the time it takes to send a frame between two locations.

Link A connection between two Fibre Channel ports consisting of a transmit fibre and a receive fibre.

Link_Control_Facility A termination card that handles the logical and physical control of the Fibre Channel link for each mode of use.

LIP A Loop Initialization Primitive sequence is a special fibre channel sequence that is used to start loop initialization. Allows ports to establish their port addresses.

Local Area Network (LAN) A network covering a relatively small geographic area (usually not larger than a floor or small building). Transmissions within a Local Area Network are mostly digital, carrying data among stations at rates usually above one megabit/s.

Login Server Entity within the Fibre Channel fabric that receives and responds to login requests.

Loop Circuit A temporary point-to-point like path that allows bi-directional communications between loop-capable ports.

Loop Topology An interconnection structure in which each point has physical links to two neighbors resulting in a closed circuit. In a loop topology, the available bandwidth is shared.

LVD Low Voltage Differential

Management Agent A process that exchanges a managed node's information with a management station.

Managed Node A managed node is a computer, a storage system, a gateway, a media device such as a switch or hub, a control instrument, a software product such as an operating system or an accounting package, or a machine on a factory floor, such as a robot.

Managed Object A variable of a managed node. This variable contains one piece of information about the node. Each node can have several objects.

Management Station A host system that runs the management software.

Meter 39.37 inches, or just slightly larger than a yard (36 inches)

Media Plural of medium. The physical environment through which transmission signals pass. Common media include copper and fiber optic cable.

Media Access Rules (MAR).

MIA Media Interface Adapter - MIAs enable optic-based adapters to interface to copper-based devices, including adapters, hubs, and switches.

MIB Management Information Block - A formal description of a set of network objects that can be managed using the Simple Network Management Protocol (SNMP). The format of the MIB is defined as part of SNMP and is a hierarchical structure of information relevant to a specific device, defined in object oriented terminology as a collection of objects, relations, and operations among objects.

Mirroring The process of writing data to two separate physical devices simultaneously.

MM Multi-Mode - See Multi-Mode Fiber

MMF See Multi-Mode Fiber - - In optical fiber technology, an optical fiber that is designed to carry multiple light rays or modes concurrently, each at a slightly different reflection angle within the optical core. Multi-Mode fiber transmission is used for relatively short distances because the modes tend to disperse over longer distances. See also: Single-Mode Fiber, SMF

Multicast Sending a copy of the same transmission from a single source device to multiple destination devices on a fabric. This includes sending to all N_Ports on a fabric (broadcast) or to only a subset of the N_Ports on a fabric (multicast).

Multi-Mode Fiber (MMF) In optical fiber technology, an optical fiber that is designed to carry multiple light rays or modes concurrently, each at a slightly different reflection angle within the optical core. Multi-Mode fiber transmission is used for relatively short distances because the modes tend to disperse over longer distances. See also: Single-Mode Fiber

Multiplex The ability to intersperse data from multiple sources and destinations onto a single transmission medium. Refers to delivering a single transmission to multiple destination Node Ports (N_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.

N_Port Login N_Port Login (PLOGI) allows two N_Ports to establish a session and exchange identities and service parameters. It is performed following completion of the fabric login process and prior to the FC-4 level operations with the destination port. N_Port Login may be either explicit or implicit.

Name Server Provides translation from a given node name to one or more associated N_Port identifiers.

NAS Network Attached Storage - a term used to describe a technology where an integrated storage system is attached to a messaging network that uses common communications protocols, such as TCP/IP.

NDMP Network Data Management Protocol

Network An aggregation of interconnected nodes, workstations, file servers, and/or peripherals, with its own protocol that supports interaction.

Network Topology Physical arrangement of nodes and interconnecting communications links in networks based on application requirements and geographical distribution of users.

NFS Network File System - A distributed file system in UNIX developed by Sun Microsystems which allows a set of computers to cooperatively access each other's files in a transparent manner.

NL_Port Node Loop Port - a node port that supports Arbitrated Loop devices.

NMS Network Management System - A system responsible for managing at least part of a network. NMSs communicate with agents to help keep track of network statistics and resources.

Node An entity with one or more N_Ports or NL_Ports.

Non-Blocking A term used to indicate that the capabilities of a switch are such that the total number of available transmission paths is equal to the number of ports. Therefore, all ports can have simultaneous access through the switch.

Non-L_Port A Node or Fabric port that is not capable of performing the Arbitrated Loop functions and protocols. N_Ports and F_Ports are not loop-capable ports.

Operation A term defined in FC-2 that refers to one of the Fibre Channel *building blocks* composed of one or more, possibly concurrent, exchanges.

Optical Disk A storage device that is written and read by laser light.

Optical Fiber A medium and the technology associated with the transmission of information as light pulses along a glass or plastic wire or fiber.

Ordered Set A Fibre Channel term referring to four 10 -bit characters (a combination of data and special characters) providing low-level link functions, such as frame demarcation and signaling between two ends of a link.

Originator A Fibre Channel term referring to the initiating device.

Out of Band Signaling Signaling that is separated from the channel carrying the information.

Peripheral Any computer device that is not part of the essential computer (the processor, memory and data paths) but is situated relatively close by. A near synonym is input/output (I/O) device.

PLDA Private Loop Direct Attach - A technical report which defines a subset of the relevant standards suitable for the operation of peripheral devices such as disks and tapes on a private loop.

PLOGI See N_Port Login

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

Port The hardware entity within a node that performs data communications over the Fibre Channel.

Port Bypass Circuit A circuit used in hubs and disk enclosures to automatically open or close the loop to add or remove nodes on the loop.

Private NL_Port An NL_Port which does not attempt login with the fabric and only communicates with other NL Ports on the same loop.

Protocol A data transmission convention encompassing timing, control, formatting and data representation.

Public NL_Port An NL_Port that attempts login with the fabric and can observe the rules of either public or private loop behavior. A public NL_Port may communicate with both private and public NL_Ports.

Quality of Service (QoS) A set of communications characteristics required by an application. Each QoS defines a specific transmission priority, level of route reliability, and security level.

RAID Redundant Array of Inexpensive or Independent Disks. A method of configuring multiple disk drives in a storage subsystem for high availability and high performance.

Raid 0 Level 0 RAID support - Striping, no redundancy

Raid 1 Level 1 RAID support - mirroring, complete redundancy

Raid 5 Level 5 RAID support, Striping with parity

Repeater A device that receives a signal on an electromagnetic or optical transmission medium, amplifies the signal, and then retransmits it along the next leg of the medium.

Responder A Fibre Channel term referring to the answering device.

Router (1) A device that can decide which of several paths network traffic will follow based on

some optimal metric. Routers forward packets from one network to another based on network-layer information. (2) A dedicated computer hardware and/or software package which manages the connection between two or more networks. See also: Bridge, Bridge/Router

SAF-TE SCSI Accessed Fault-Tolerant Enclosures

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.

SAN System Area Network - term originally used to describe a particular symmetric multiprocessing (SMP) architecture in which a switched interconnect is used in place of a shared bus. Server Area Network - refers to a switched interconnect between multiple SMPs.

SC Connector A fiber optic connector standardized by ANSI TIA/EIA-568A for use in structured wiring installations.

Scalability The ability of a computer application or product (hardware or software) to continue to function well as it (or its context) is changed in size or volume. For example, the ability to retain performance levels when adding additional processors, memory and/or storage.

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. The table below identifies the major characteristics of the different SCSI version.

SCSI Version	Signal Rate MHz	Bus-Width (bits)	Max. DTR (MB/s)	Max. Num. Devices	Max. Cable Length (m)
SCSI-1	5	8	5	7	6
SCSI-2	5	8	5	7	6
Wide SCSI-2	5	16	10	15	6
Fast SCSI-2	10	8	10	7	6
Fast Wide SCSI-2	10	16	20	15	6
Ultra SCSI	20	8	20	7	1.5
Ultra SCSI-2	20	16	40	7	12
Ultra 2 LVD SCSI	40	16	80	15	12

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).

SCSI-FCP The term used to refer to the ANSI Fibre Channel Protocol for SCSI document (X3.269-199x) that describes the FC-4 protocol mappings and the definition of how the SCSI protocol and command set are transported using a Fibre Channel interface.

Sequence A series of frames strung together in numbered order which can be transmitted over a Fibre Channel connection as a single operation. See also: Exchange

SERDES Serializer Deserializer

Server A computer which is dedicated to one task.

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.

Single-Mode Fiber In optical fiber technology, an optical fiber that is designed for the transmission of a single ray or mode of light as a carrier. It is a

single light path used for long-distance signal transmission. See also: Multi-Mode Fiber

SMART Self Monitoring and Reporting Technology

SM Single Mode - See Single-Mode Fiber

SMF Single-Mode Fiber - In optical fiber technology, an optical fiber that is designed for the transmission of a single ray or mode of light as a carrier. It is a single light path used for long-distance signal transmission. See also: MMF

SNIA Storage Networking Industry Association. A non-profit organization comprised of more than 77 companies and individuals in the storage industry.

SN Storage Network. See also: SAN

SNMP Simple Network Management Protocol - The Internet network management protocol which provides a means to monitor and set network configuration and run-time parameters.

SNMWG Storage Network Management Working Group is chartered to identify, define and support open standards needed to address the increased management requirements imposed by storage area network environments.

SSA Serial Storage Architecture - A high speed serial loop-based interface developed as a high speed point-to-point connection for peripherals, particularly high speed storage arrays, RAID and CD-ROM storage by IBM.

Star The physical configuration used with hubs in which each user is connected by communications links radiating out of a central hub that handles all communications.

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 (e.g. 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 including -- reporting of capacity, performance, etc. 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 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.

Striping A method for achieving higher bandwidth using multiple N_Ports in parallel to transmit a single information unit across multiple levels.

STP Shielded Twisted Pair

Storage Media The physical device itself, onto which data is recorded. Magnetic tape, optical disks, floppy disks are all storage media.

Switch A component with multiple entry/exit points (ports) that provides dynamic connection between any two of these points.

Switch Topology An interconnection structure in which any entry point can be dynamically connected to any exit point. In a switch topology, the available bandwidth is scalable.

T11 A technical committee of the National Committee for Information Technology Standards, titled T11 I/O Interfaces. It is tasked with developing standards for moving data in and out of computers.

Tape Backup Making magnetic tape copies of hard disk and optical disc files for disaster recovery.

Tape Pooling A SAN solution in which tape resources are pooled and shared across multiple hosts rather than being dedicated to a specific host.

TCP Transmission Control Protocol - a reliable, full duplex, connection-oriented end-to-end transport protocol running on top of IP.

TCP/IP Transmission Control Protocol/ Internet Protocol - a set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks.

Time Server A Fibre Channel-defined service function that allows for the management of all timers used within a Fibre Channel system.

Topology An interconnection scheme that allows multiple Fibre Channel ports to communicate. For example, point-to-point, Arbitrated Loop, and switched fabric are all Fibre Channel topologies.

Twinax A transmission media (cable) consisting of two insulated central conducting leads of coaxial cable.

Twisted Pair A transmission media (cable) consisting of two insulated copper wires twisted around each other to reduce the induction (thus interference) from one wire to another. The twists, or lays, are varied in length to reduce the potential for signal interference between pairs. Several sets of twisted pair wires may be enclosed in a single cable. This is the most common type of transmission media.

ULP Upper Level Protocols

UTC Under-The-Covers, a term used to characterize a subsystem in which a small number of hard drives are mounted inside a higher function unit. The power and cooling are obtained from the system unit. Connection is by parallel copper ribbon cable or pluggable backplane, using IDE or SCSI protocols.

UTP Unshielded Twisted Pair

Virtual Circuit A unidirectional path between two communicating N_Ports that permits fractional bandwidth.

WAN Wide Area Network - A network which encompasses inter-connectivity between devices over a wide geographic area. A wide area network may be privately owned or rented, but the term usually connotes the inclusion of public (shared) networks.

WDM Wave Division Multiplexing - A technology that puts data from different sources together on an optical fiber, with each signal carried on its own separate light wavelength. Using WDM, up to 80 (and theoretically more) separate wavelengths or channels of data can be multiplexed into a stream of light transmitted on a single optical fiber.

WEBM Web-Based Enterprise Management - A consortium working on the development of a series of standards to enable active management and monitoring of network-based elements.

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.

Index

D

E

F

G

H

I

P

Q

R

S

T

U

V

W

Z

IBM Redbooks review

Your feedback is valued by the Redbook authors. In particular we are interested in situations where a Redbook "made the difference" in a task or problem you encountered. Using one of the following methods, **please review the Redbook, addressing value, subject matter, structure, depth and quality as appropriate.**

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

Document Number **Redbook Title**	SG24-6116-00 Planning and Implementing an IBM SAN
Review	
What other subjects would you like to see IBM Redbooks address?	
Please rate your overall satisfaction:	O Very Good O Good O Average O Poor
Please identify yourself as belonging to one of the following groups:	O Customer O Business Partner O Solution Developer O IBM, Lotus or Tivoli Employee O None of the above
Your email address: The data you provide here may be used to provide you with information from IBM or our business partners about our products, services or activities.	O Please do not use the information collected here for future marketing or promotional contacts or other communications beyond the scope of this transaction.
Questions about IBM's privacy policy?	The following link explains how we protect your personal information. **ibm.com**/privacy/yourprivacy/

Planning and Implementing an IBM SAN

(0.5” spine)
0.475”<->0.875”
250 <-> 459 pages

Planning and Implementing an IBM SAN

Review planning considerations associated with Fibre Channel products

Learn how to build an IBM SAN environment

Set up and administer zoning

This book is a follow-on from the IBM Redbook, *Designing an IBM Storage Area Network,* SG24-5758 . In that book we introduced Fibre Channel basics, described the technical topology of a SAN, detailed Fibre Channel products, and IBM SAN initiatives. We also designed configurations that were able to maximize the benefits of Fibre Channel products that are currently supported by IBM, and available in the marketplace today.

Where this IBM Redbook picks up the story is how to implement those products that are in the IBM product armory today. It is not possible to duplicate each and every SAN installation that is possible, feasible, or practical. What we want to achieve is a consolidated reference guide that details how the basic products can be swiftly and, in some cases, easily implemented. We will show the various features that each of these products enjoys, and how the most common and important benefits of each are taken advantage of. We will show how they can be employed in some of the more commonly encountered environments and platforms.

With this in mind, in this redbook we have two objectives. The first is to show practical decisions to be considered when planning a SAN; the second objective is to show how the following products can be installed, configured and tailored:

IBM Enterprise Storage Server with native Fibre Channel
IBM Fibre Channel Storage Hub
IBM SAN Fibre Channel Managed Hub
IBM SAN Fibre Channel Switch
McDATA Enterprise Fibre Channel Director
IBM Storage Area Network Data Gateway
Vicom Fibre Channel SLIC Router

Once these products are successfully installed, and all these configurations have been tested using a "hands-on" environment, we will show some of the benefits that are fundamental to their application in a SAN.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6116-00 ISBN 073841829X