IBM

Implementing Fibre Channel Attachment on the ESS

Fast connectivity for a fast disk subsystem

Steps to migrate from SCSI connectivity to Fibre Channel

Positions the ESS at the heart of your SAN

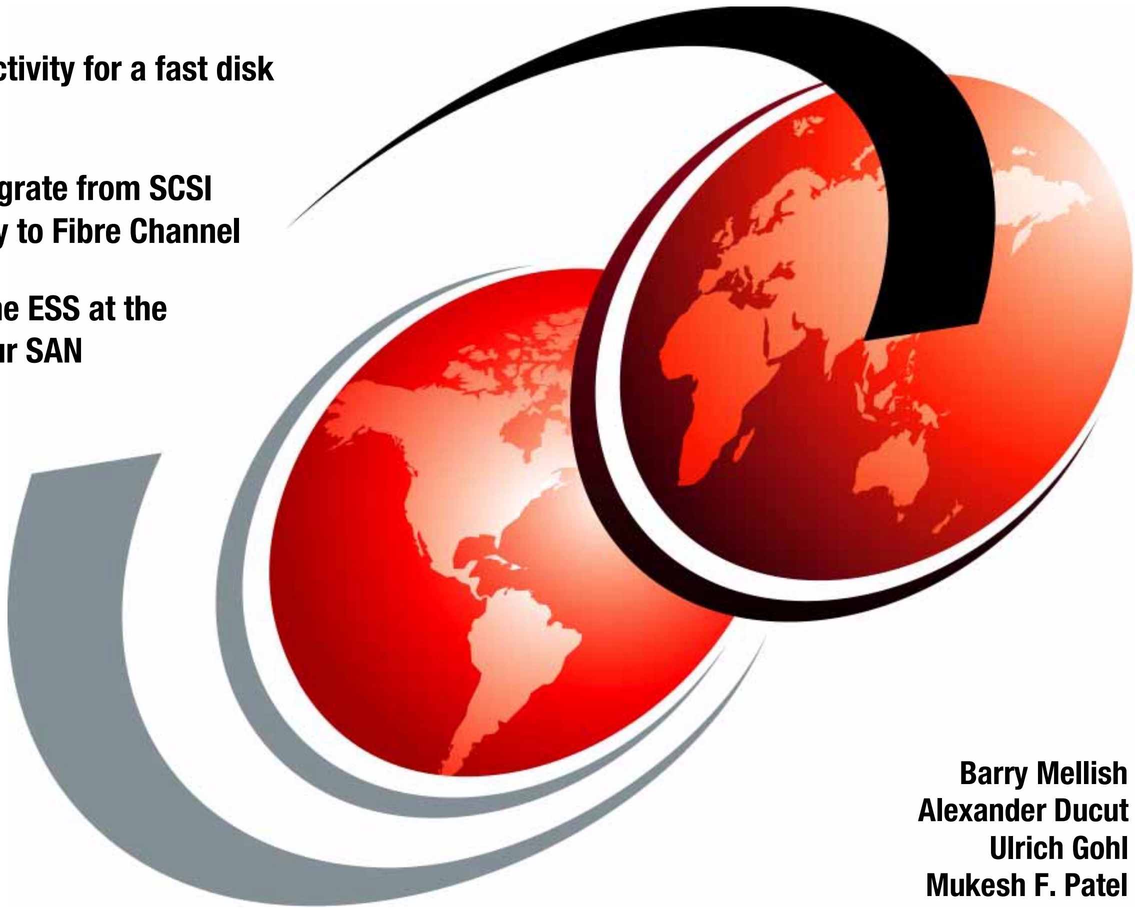

Barry Mellish
Alexander Ducut
Ulrich Gohl
Mukesh F. Patel

Redbooks

ibm.com/redbooks

SG24-6113-00

International Technical Support Organization

Implementing Fibre Channel Attachment on the ESS

September 2000

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix D, "Special notices" on page 109.

First Edition (September 2000)

This edition applies to the Enterprise Storage Server models 2105-Exx and 2105-Fxx.

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

Contents

Figures

Tables

Preface

This IBM Redbook will help you install, tailor, and configure Fibre Channel (FC) attachment of Open Systems hosts to the IBM Enterprise Storage Server (ESS).

This redbook gives a broad understanding of the procedures involved and describes the prerequisites and requirements. It then shows you how to implement FC attachment for each of the announced Open Systems platforms. The book also describes the steps required to migrate to direct FC attachment from native SCSI adapters and from FC attachment via the SAN Data Gateway (SDG).

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.

Barry Mellish is a Project Leader at the International Technical Support Organization, San Jose Center. He has coauthored three previous redbooks and has taught many classes on storage subsystems. He joined IBM UK sixteen years ago, and before recently joining the ITSO, he worked as a senior storage specialist in the Disk Expert team in EMEA.

Alexander Ducut is working in PC Technical Support at IBM Philippines. He has six years of experience in the Information Technology field. His areas of expertise include PC hardware, network operating systems, storage hardware and software. He conducts several PC related courses for service centers and business partners, and has coauthored the redbook *Using LCCM Functions with Servers and Workstations*. He is currently an MCSE and PSE.

Ulrich Gohl is a Senior HW Specialist. He joined IBM in 1978 as a Hardware Engineer. He has practical experience with ESCON problem determination at many OS/390 customers in Germany. Ulrich's areas of expertise include all types of S/390 attachable disk subsystems as well as the S/390 servers. He has considerable practical experience in ESS implementations, including S/390 and Open systems.

Mukesh F. Patel is an IT Architect developing storage management solutions for IGS Service Delivery Center South. He has 15 years of experience in the Information Technology field. He holds a master's degree in Computer Science from CUNY. His areas of expertise include storage management hardware, software, UNIX, and systems management. He joined IBM three years ago. He worked at CSC Consulting and Systems Integration prior to joining IBM. He has developed numerous storage management solutions at IBM.

The team would also like to extend their thanks to the following people for their invaluable contributions to this project:

- Jack Flynn, Device Technical Office, VM/VSE/Open Systems, IBM San Jose
- Phil Mills, DASD Subsystems Architect, IBM San Jose
- Tom Connolly, IBM San Jose
- Paula Collins Di Benedetto, Shark SLS Team Leader, IBM San Jose
- Brian Smith, IO Subsystem Performance Evaluation, IBM San Jose
- Neena Cherian, IO Subsystem Performance Evaluation, IBM San Jose
- John Aschoff, IO Subsystem Performance Evaluation, IBM San Jose
- Bill Micka, Subsystem Architecture, IBM Tucson
- Dave Reeve, Storage Consultant, IBM UK
- Paul Hurley, DST test, IBM San Jose
- Tammy Sokol, Manager EVT Group and Manager ESS EVT Testing, IBM San Jose
- Ronda Hruby, EVT Team Leader, IBM San Jose
- John Young, EVT Test Team, IBM San Jose
- Ann Rudy, Shark Program Management, IBM San Jose
- Ruth Azevedo, ESS Specialist Development Team, IBM San Jose
- Alison Pate, Technical Support for SSD Software Products, IBM San Jose
- Pat Blaney, Technical Support for SSD Software Products, IBM San Jose
- Ling Pong, Technical Support for SSD Software Products, IBM San Jose
- Edsel Carson, Information Development, IBM San Jose
- Henry Caudillo, Information Development, IBM San Jose
- Marilyn Pullen, Information Development, IBM San Jose
- Rainer Wolafka, EMEA Storage Center, IBM Mainz
- Jon Tate, ITSO SAN Solutions Project Leader, IBM San Jose
- Mark Blunden, ITSO Open Storage Project Leader, IBM San Jose
- John Kotches, USF Development, IBM Schaumburg

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 135 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

The evolution of Information Technology (IT), coupled with the explosion of the Internet has led to a new business model that is commonly referred to as e-business or e-commerce.

E-business is a strategic imperative for many companies worldwide. The driving force behind this is the need to cut costs, thereby improving profits. A second but equally important reason is the need and desire to expand a company's reach. e-business allows you, as a company, to communicate easily with your customers anytime, anywhere.

E-business is network-centric, global and data-intensive. The most widely used network today is the Internet. The Internet is global, it can be accessed from anywhere around the globe.

Today we are faced with an explosion of data generated by e-business applications. This data must be stored and managed appropriately. These trends are driving the requirement for more flexible and reliable access to data. The data resides on storage devices. This is driving the requirement for a Storage Area Network (SAN).

Beside the previous mentioned trends, heterogeneous IT environments are very common. Today's state-of-the-art storage resources must be able to support whatever platforms are within the environment. This is essentially an investment protection requirement that allows you to configure a storage resource for one set of systems and subsequently configure part of the capacity to other systems on an as-needed basis.

The IBM Enterprise Storage Server (ESS) is one of the most flexible, powerful, scalable and reliable disk subsystems available on the marketplace. The ESS can be connected universally by using Enterprise System Connectivity (ESCON) for S/390 attachment and SCSI for Open Systems attachment.

Today, many AIX, Windows NT and other Open Systems based applications have rapidly growing storage requirements in terms of capacity and connectivity.

The capacity issue has been covered already by ESS itself. The ESS supports concurrent upgrades of storage capacity as well as the dynamic assignment of this storage to a mix of heterogeneous servers. Concurrent in this case means concurrent from ESS perspective. If the attached servers operating system (OS) does not have a function that allows disk drives to be added “in flight”, this server may need to be rebooted. This is not concurrent for the server that reboots, but for others that may be attached to the same ESS.

The new Fibre Channel (FC) attachment on the ESS addresses the connectivity issue as well. The limitations seen on SCSI in terms of distance, addressability and performance are gone with Fibre Channel attachment.

Fibre Channel attachment is not just a replacement of the parallel SCSI interface based on copper cables by a serial, fiber optic based interface. In view of today’s operating system islands, it is more likely entry into future universal data access. The keyword for that is SAN.

Fibre Channel is the IT industry’s strategic new interface that will allow customers the most flexibility in achieving future goals in storage growth. It is so popular, even today, that every major storage vendor is designing and shipping products that are Fibre Channel capable. This includes storage subsystems, hosts, switches, hubs, repeaters and especially the software to manage the various Fibre Channel configurations.

The Fibre Channel technology promises to avoid any early death due to obsolescence by being built upon a scalable architecture that inherently supports signaling rate improvements of at least one order of magnitude (to 10 Gbit/sec) and probably greater. This scalability was a major objective of the Fibre Channel architecture from its very beginning.

It is the very rich connectivity options made possible by Fibre Channel technology that has fueled the concept of SANs in the last year or two. Fibre Channel is the infrastructure of the SAN.

The ESS can be implemented into a SAN by using this new attachment option. Being a member of a SAN enables the ESS to fully participate on future SAN applications, like LAN-less or LAN-free backup and others. For more detailed information refer to *Introduction to Storage Area Network,SAN*, SG24-5470. The best is yet to come.

Fiber optic technology is a well known and proved technology within IBM. It has been used since 1991, when IBM introduced ESCON on S/390 based host systems. This laid the groundwork for SAN. ESCON, the industry’s first SAN, started out as a serial, fiber optic, point-to-point switched network connecting disk, tape and printing devices to S/390 (mainframe) servers.

IBM has the knowledge to provide many service offerings to help customers to build up a fiber optic infrastructure in their Data Processing Center and to design and implement a SAN.

This Redbook is intended to help IBM customers by providing information on what Fibre Channel is and how to implement the new ESS Fibre Channel attachment into their environment. In-depth information about how to design and implement a SAN can be found in the redbooks *Designing an IBM Storage Area Network*, SG24-5758 and *Planning and Implementing an IBM SAN*, SG24-6116.

It also enables IBM customers to plan the steps needed for a migration from SCSI to Fibre Channel attachment, whether native or by using a SAN Data Gateway (SDG).

Chapter 2. Fibre Channel basics

The purpose of this chapter is to describe what Fibre Channel (FC) is all about. The intent is to give the IBM Enterprise Storage Server (ESS) users who are responsible for planning and configuring storage an idea how they should best plan for using this new attachment in their environment.

2.1 Introduction

Fibre Channel is a technology standard that allows data to be transferred from one node to another at high speeds. This standard has been defined by a consortium of industry vendors and has been accredited by the American National Standards Institute (ANSI).

Fibre Channel standard has five layers. They are defined as FC-0 through FC-4. These layers are described in detail in the primary Fibre Channel Specification. FC-0 through FC-2 define how Fibre Channel ports interact through their physical links to communicate with other ports. FC-3 through FC-4 define how Fibre Channel ports interact with applications in host systems.

The word ***FIBRE*** in Fibre Channel takes the French spelling rather than the traditional spelling of fiber, as in fiber optics. This is because the interconnection between nodes are not necessarily based on fiber optics, but can also be based on copper cables.

Note

Fibre Channel on ESS is always based on fiber optics.

When we talk about Fibre Channel connections within this redbook, we always mean fiber optics rather than copper cables.

This new attachment feature on ESS provides the capability to build a Storage Area Network (SAN) that is able to interconnect host systems and storage devices which are far away from each other. The distance limits are improved dramatically compared to copper-cable limitations on SCSI.

Fibre Channel is an architecture used to carry traffic using various protocols. Fibre Channel Protocol (FCP) is the one used to carry SCSI traffic. FCP is expected to become the standard protocol for Open Systems, while FICON will become the standard protocol for S/390 systems. FC can also carry data traffic that uses IP protocol.

Fibre Channel is a 100 MB/sec, full duplex, serial communications technology. Depending on the fiber optics used, the physical fiber interface used to interconnect two node ports, called a link, can be up to ten kilometers in length. Fibre Channel merges the best features of networking interfaces, like connectivity and scalability with the best features of traditional interfaces, like reliability and throughput.

A SAN, built upon Fibre Channel, consists of elements like switches, hubs, and bridges that are used to interconnect Fibre Channel ports.

2.2 Fiber optic technology

Fiber optic technology uses light pulses traveling through a *cable* made of silica glass, called an optical fiber. These light pulses transport data-signals through the fiber. It can be compared to the way electrical pulses travel through a copper-cable transporting data-signals.

The advantages of optical fiber are:

- High bandwidth
- No electromagnetic interference
- Very secure
- All dielectric — eliminates grounding and ground potential problems
- Less weight compared to copper-cables

2.2.1 How do fiber optics function?

An optical fiber has a central core and a surrounding cladding of slightly different glass material. If a light source is used to put light into the fiber, these two different glass materials are working similar to a mirror, reflecting the light.

The physical size of an optical fiber is determined by the diameter of the core and cladding, measured in microns. One micron is a millionth (1/1,000,000) of a meter. An optical fiber cable having a core diameter of 50 microns and a cladding of 125 microns, is designated as 50/125 microns. Other fibers commonly used are 62.5/125 and 9/125. The most common one used for Fibre Channel adapters will probably be the 50/125 microns fiber.

Figure 1 illustrates how the light is propagated through the different fibers.

.

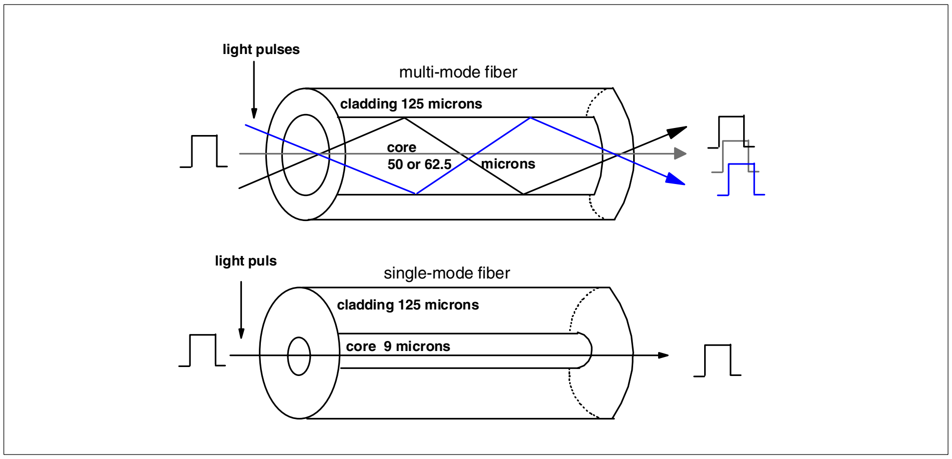

Figure 1. Light propagation through multi-mode and single-mode fiber

Fiber optic connections within a SAN are built upon two fibers. Each port, whether it is on a host, a storage-unit, a hub or a switch does have a receiver site and a transmitter site. To ensure proper communication a transmitter on one port is always connected via a fiber to a receiver on another port and vice versa. Figure 2 shows an example.

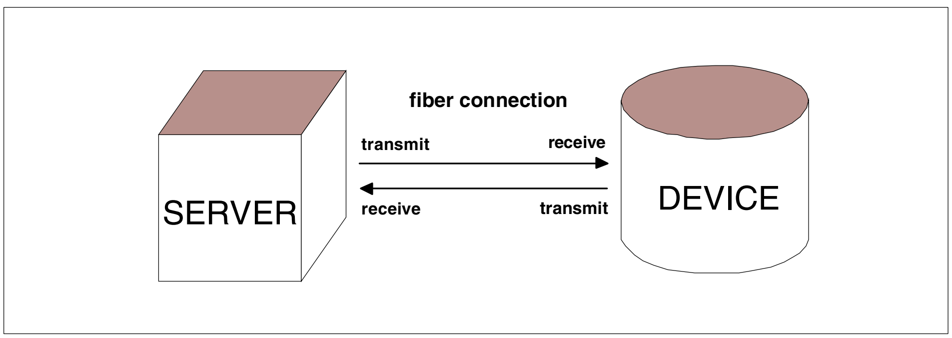

Figure 2. An example of fiber connection

Depending on the topology you choose, the cabling infrastructure may vary. See 2.5, “Fibre Channel topologies” on page 13 for more information on topologies.

Different laser technologies can be used to transmit the light pulses. Depending on the laser and fiber used, different link lengths are supported. It is necessary that the fiber optic cables used match the requirements.

Table 1 shows the various combinations of lasers and fibers as well as the resulting link length. As a base rule, you should use single-mode fibers for long-wave laser and multi-mode fibers for short-wave laser. If the infrastructure of your DP-Center already has multi-mode fibers available, they can be used to interconnect long-wave laser equipped ports, but in that case a so called mode-conditioner is required. For certain, the link length is limited to 175meters/500meters accordingly.

Sections 2.2.2, “Short wave laser” on page 8 and 2.2.3, “Long wave laser” on page 9 provide further information.

Table 1. Link lengths

Fiber size	Wave length	Fiber mode	Link length
9 / 125 microns	1300nm, long-wave	single	10 km
50 / 125 microns	1300nm, long-wave	multi *	500 meters
62.5 / 125 microns	1300nm, long-wave	multi *	175 meters
50 / 125 microns	780nm, short-wave	multi	500 meters
62.5 / 125 microns	780nm, short-wave	multi	175 meters

* A mode-conditioner is required for such a link.

Two different laser optical transmitter technologies are supported by the Fibre Channel architecture. They are:

- Short wave laser
- Long wave laser

Note

The ESS uses short wave laser. Therefore the first hop from the ESS has a maximum distance of 500 meters. Subsequent hops can be longer.

2.2.2 Short wave laser

The short wave laser uses a wavelength of 780 nanometers. This laser is compatible with multi-mode fiber, only. A maximum distance of 500 meters between two ports is supported with a 50/125 micron fiber. A 62.5/125 micron fiber lowers the supported distance to 175 meters.

2.2.3 Long wave laser

The long wave laser uses a wavelength of 1300 nanometers. Long wave lasers are compatible with both single-mode and multi-mode fiber. The maximum distance between two ports (10 kilometers) can be achieved with a 9/125 micron single-mode fiber, only. A so called mode-conditioner is required to use multi-mode fibers together with long wave laser. In case that a multi-mode fiber is used to interconnect Fibre Channel adapters equipped with long wave laser, the same distance limits apply as for short wave laser, 500 meters with a 50/125 micron fiber and 175 meters with a 62.5/125 micron fiber.

2.3 Fiber infrastructure

SCSI interconnections between adapters are mostly built up by one cable assembly. Due to the length limitation of 25 meters, there was no need for a cabling infrastructure.

Fibre Channel connections do have larger length limitations, so probably they will in many cases not consist of one optical fiber cable.

Different types of fiber optic cables are available. Depending on the environmental condition and the type of installation, you choose the fiber optic cable according to your requirements. The different fiber cable types are:

- Jumper cables
- Multi-jumper cables
- Trunk cables

2.3.1 Jumper cables

A optical fiber cable, physically is a pair of optical fibers that provide two unidirectional serial bit transmission lines. Such a fiber cable is commonly referred to as a jumper cable.

Jumper cables are not protected against mechanical influences. They are commonly used for short fiber optical links between servers and storage devices within the same room.

.

Note

IBM will supply one 50/125 microns jumper cable with a length of 31 meters for each Feature Code 3022 (Fibre Channel attachment on ESS) ordered.

2.3.2 Multi-jumper cables

There are fiber cables available with more than one pair of fibers. They are called multi-jumpers. They are commonly used for fiber optical links within the same building. Multi-jumpers are used in large DP-Centers to interconnect more than one pair of Fibre Channel ports.

2.3.3 Trunk cables

Trunk cables typically contain 12 to 144 fibers. They have a strength member and an outer jacket. Each Fibre Channel link requires two trunk fibers. The physical trunk cable configuration varies and depends on the user requirements, environmental conditions, and the type of installation, for example, aboveground or underground.

2.3.4 Fibre Channel link configuration

To gain, for example a distance of 175 meters between two buildings, you most likely will see a infrastructure built up with multiple fiber cables.

A Fibre Channel optical link configuration schema may be similar to the one shown in Figure 3.

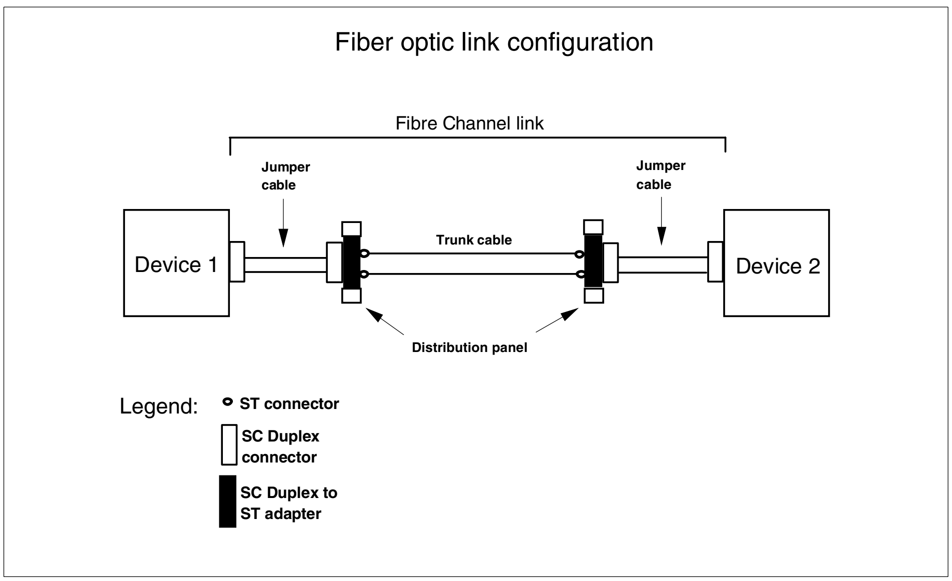

Figure 3. Fiber optic link configuration

It may be a good idea to have a distribution- or patch-panel on every floor. All Fibre Channel ports which require a connection to a port on another floor or building can be connected to the floor distribution- / patch-panel by using jumper-cables. The floor distribution- / patch-panels are interconnected by using multi-jumpercables, while the buildings are interconnected using a trunk-cable.

A *SC Duplex Connector* is a fiber optic connector standardized by ANSI TIA/EIA-568A for use in structured wiring installations.

Figure 4 illustrates how an infrastructure may look.

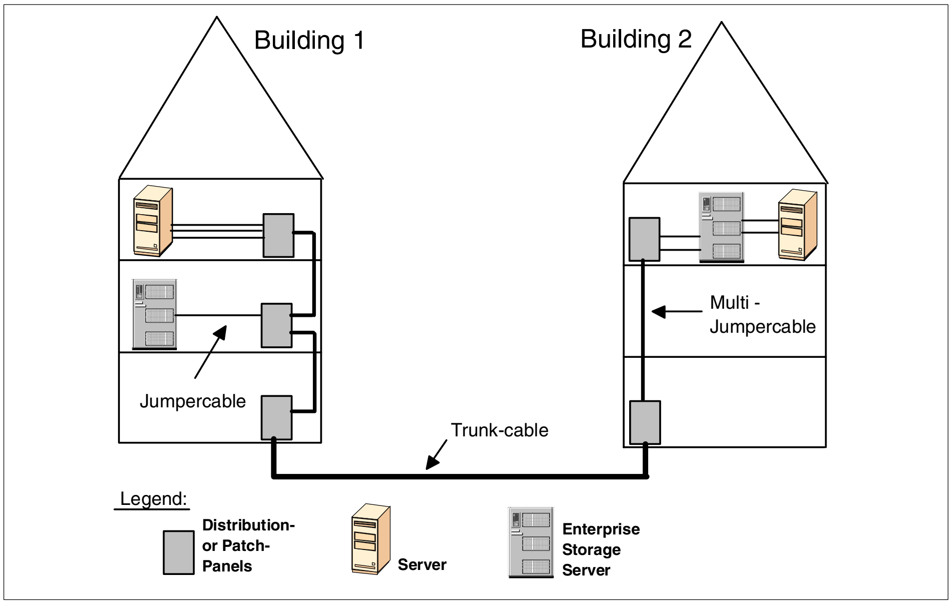

Figure 4. An example of fiber cabling infrastructure

A infrastructure like the previous one allows you to reconfigure your systems without the need for recabling between the floors. You just need to rearrange your links within the distribution- / patch-panels.

The following URL can be used for additional information about the connectivity services that IBM Global Services can offer you:

`http://www.as.ibm.com/asus/connectivity.html`

2.3.5 Light Loss Budget

A typical fibre optic cable installation such as that in Figure 3 will not be a continuous run of fiber optic cable. The main trunk cable may have several splices or joints and the patch panels themselves will introduce additional connections into the cable run. Every time a joint is made in a fiber optic cable there will be a slight loss of light, with the patch panel cable connections themselves introducing considerably more light loss.

When designing a fiber cable infrastructure it is important not to exceed the light loss budget. For fiber channel the typical light loss budget for a short wave link is 8db, which means that the amount of light lost through all joints and connections must not exceed 8db. It is possible to calculate the light loss budget in a fiber optic link before it is made, as there are standard light loss figures available for each type of joint or connection. Fibre Optic cable installers also have test equipment that can measure the light loss down a link as well as the actual distance of the link.

2.3.5.1 Connection of devices

In many customers the ESS, the Servers and the switches will be in the same computer room, with simple fiber connections directly from the ESS port to the switch port and from the switch port direct to the server port. In this situation the fiber cable installation work may be done "in-house" by the customer. The cables used for these connections may well be of the standard twin wire type that ships with the ESS FC ports. This type of cable is quite fragile and care needs to be taken with its installation, for example, it must not be bent in too tight a radius. It is also easily damaged by heavy objects, so it is best not laid carelessly on the underfloor. For a reliable cable installation it is recommended that cables are installed by a competent cable installer.

2.3.5.2 Patch Panels

It is common practice to use patch panels in fiber optic installations, particularly when connecting between floors or buildings. The theory is that this creates a flexible infrastructure, but in practice most customers do not change fiber connections around very often. Care needs to be taken with patch panels because the connectors cause the highest amount of light loss in the link. A speck of dirt can cause several db of light loss which could render a link inoperative. Connectors should always be cleaned whenever they are replugged.

For long haul links, such as those used for PPRC, it is worth considering making the whole link between the ESS's one continuous run of cable with spliced joints instead of connectors. This will minimize the light loss and makes the installation tamper-proof.

2.4 Nodes and ports

Each host or storage device connected to a SAN is called a node and has a World Wide Node Name (WWNN) assigned to it.

A node can have multiple FC adapters. Most FC adapters have one fiber connector, called a port. Every port, whether it is on a host, a storage device, a switch or a SAN Data Gateway, has a World Wide Port Name (WWPN) assigned to it. This WWPN is unique. If a FC adapter has two ports, it will have two WWPNs, too.

2.5 Fibre Channel topologies

Three different Fibre Channel topologies are defined in the Fibre Channel architecture. All three are supported by ESS. These are:

- Point-to-point
- Switched fabric - FC-SW
- Arbitrated loop - FC-AL

2.5.1 Point-to-point

The simplest topology is a point-to-point bi-directional connection between two nodes. By using a fiber cable, two Fibre Channel adapters (one host and one storage device) are interconnected. This topology satisfies only basic requirements. It supports the maximum bandwidth capability of Fibre Channel. It does not provide sufficient connectivity to support complex configurations. It is used when future expansion is not predicted. This topology offers limited improvements versus SCSI, but it does not exploit all the benefits a SAN may have.

ESS fully supports the point-to-point topology.

Figure 5 illustrates a point-to-point topology. Because the ESS supports only short wave laser, a multi-mode fiber has to be used. The distance between the server and the ESS can be up to 500 meters if a 50/125 micron fiber is used.

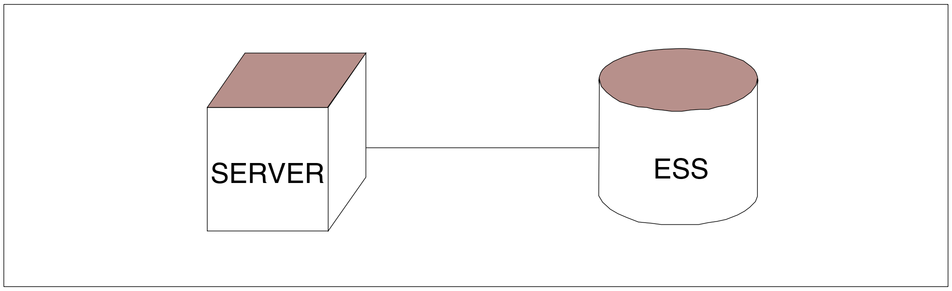

Figure 5. An example of point-to-point topology

Ports connected using this topology are designated as N_ports.

2.5.2 Switched fabric

Whenever a switch is used to interconnect Fibre Channel adapters, we talk about a switched fabric. A switched fabric is an intelligent switching infrastructure, which delivers data from any source to any destination. Figure 6 shows an example of a switched fabric.

Functions on switches and storage devices, like zoning and LUN masking, are available to provide users with the possibility of preventing servers from accessing storage they are not allowed to access. How to setup LUN masking on the ESS is described in 3.1.1, “LUN masking” on page 19.

A fabric can be built with just one switch. The Fibre Channel architecture also allows multiple switches to be interconnected, this is called switch cascading. When switches are cascaded, multiple paths between a host system and the storage device will be created. However, it is important to note that only the shortest paths between switches will be used. Both the Brocade and McData switches use a protocol called "Find Shortest Path First" (known as FSPF). This calculates the shortest path between a host and a device during fabric initialization and from then on, only that path and others equal in length to it will be used. If there are alternative routes through other switches which are longer they will not be used until all the shortest paths have failed.

The use of inter-switch links therefore needs careful consideration. Links between switches are desirable so that a common zoning configuration can exist across the fabric; it also means switches can be managed using IP over Fibre from only one switch ethernet connection.

Multiple links between switches only make sense if all the route lengths between host and storage are the same.

Most switches available at present support 8, 16, 32 or 64 ports.

Ports used to interconnect switches are designated as E_ports. Ports connected to the fabric coming from a server or a storage device can be N_ports

The Fibre Channel architecture allows to address up to 256 switches within a fabric. Because some of the possible addresses are reserved, there are only 239 addresses actually available. This means you can have as many as 239 switches within one switched fabric SAN environment.

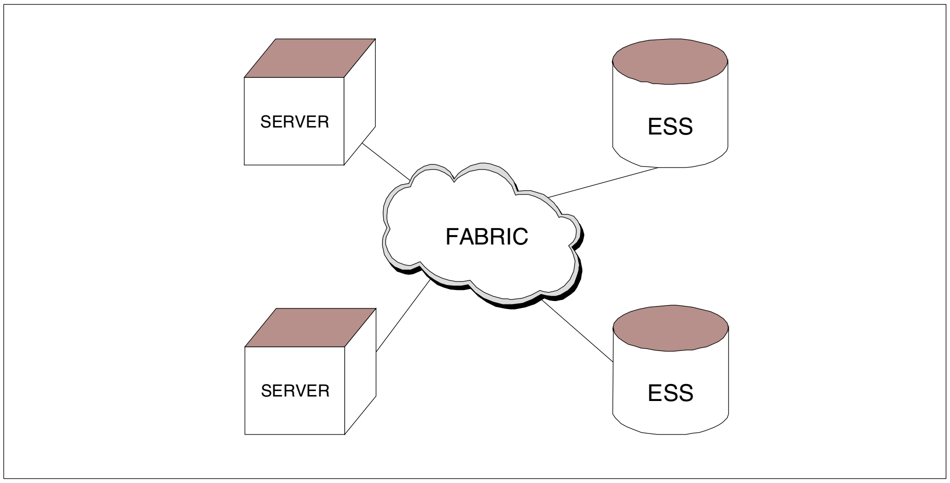

Figure 6. An example of switched fabric

The switched fabric topology provides you with all the benefits Fibre Channel offers:

- Flexibility
- Scalability

2.5.3 Arbitrated loop

Arbitrated loop topology is a unidirectional ring topology, similar to token ring. Up to a maximum of 127 Fibre Channel ports can be interconnected via a looped interface. Information is routed around the loop and repeated by intermediate ports until it arrives at its destination. See Figure 7 for an example of a simple arbitrated loop. The Fibre Channel ports that support this topology must be able to perform these routing and repeating functions in addition to all the functions required by the point-to-point ports. The acronym FC-AL often refers to this topology. All ports share the FC-AL interface and therefore also share the bandwidth of the interface. Only one connection may be ongoing at a time. Ports connected to FC-AL are called NL_ports.

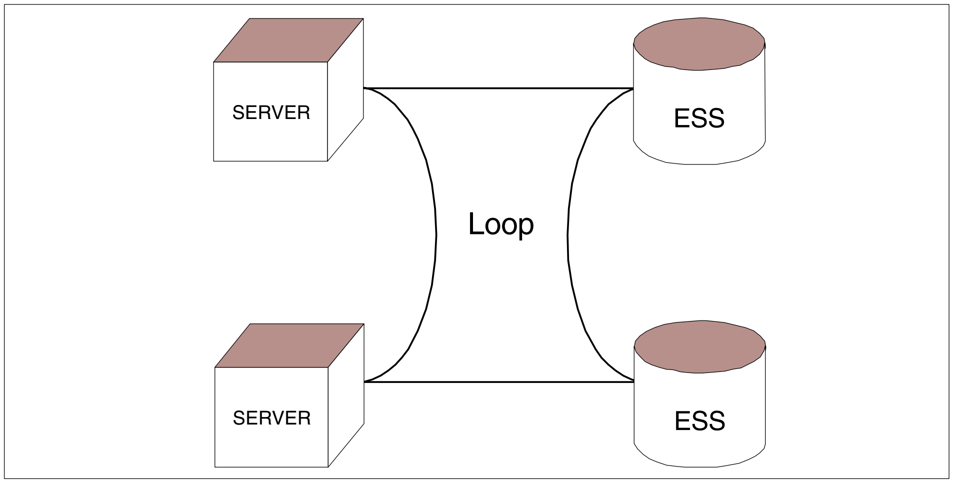

Figure 7. An example of arbitrated loop

A very commonly used variation on this topology is the Fibre Channel hub, shown in Figure 8.

A hub incorporates the structure of FC-AL in a package that provides ports physically connected similar to those of a switch. This allows more centralized cabling and, due to added functionality provided by hubs, increases the reliability of the arbitrated loop topology. The hub operates in a similar way to a multiple access unit (MAU) in a token ring. It has the ability to switch out a failing component so that the integrity of the ring is not compromised. It does not, however, improve the performance of the loop, since it is still just a single loop that has been repackaged.

Today's hubs commonly have 8, 16, or 32 ports. Additional ports can be configured by cascading hubs. This cascading is done by interconnecting hubs. Remember, the resulting configuration is still a single loop.

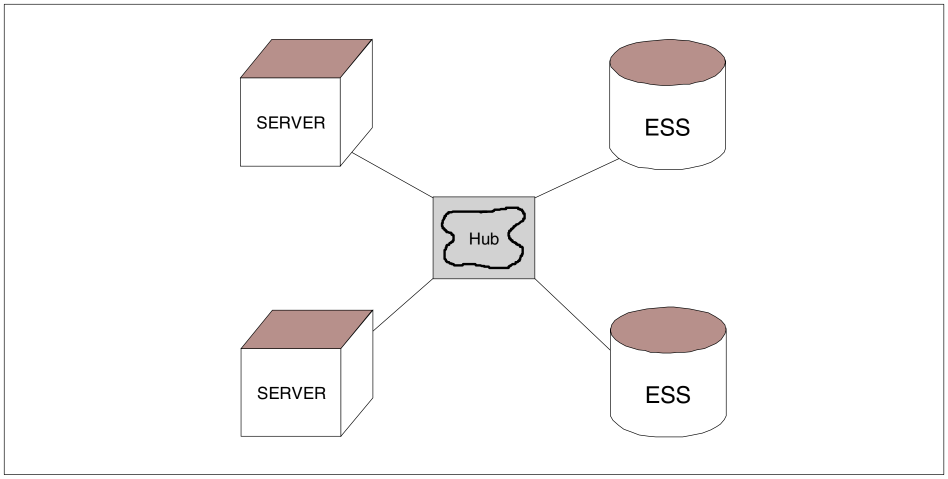

Figure 8. An example of arbitrated loop with a hub

Fiber optics infrastructure for a FC-AL with a hub is the same as a switched fabric infrastructure.

2.6 Interoperability

There are many vendors that have products that form a storage area network (SAN), for example, Host Bus Adapter, switches, storage products and others. Due to slightly different implementations of the Fibre Channel standard by the various vendors and the wide range of configuration options, the topic interoperability has to be considered when planning a SAN. Even if you can build a fabric out of multiple switches, it may not be possible to build a fabric out of switches from different vendors. To help customers resolve these issues, IBM and other vendors have built SAN Interoperability LABs. These LABs are proving grounds where the interoperability of different hardware and software in a SAN environment can be tested. They can also carry out a proof of concept for your design, by testing it at their location instead of testing it in your production environment.

The fee for such testing may varies depending on your requirements. IBM provides information about which products and configurations are supported by IBM. You should follow these guidelines at all times. As an example, the following URLs provide information on the IBM2109 switches and IBM105 Enterprise Storage Server:

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

Chapter 3. ESS Specialist

The IBM ESS comes with a network based tool that allows managing and monitoring of the ESS. This product is called ESS Specialist. It is a Web based utility for configuration and administration, using a secure Internet connection normally using a PC. You need a browser (such as Netscape Navigator, Microsoft Internet Explorer, or Sun Hot Java) that supports Java 1.1.6 or higher.

The ESS Specialist provides the customer with the ability to:

- Monitor error logs
- View the status of the ESS
- View and update the configuration
- Add, delete or modify host systems
- Configure host ports (both SCSI and fibre)
- Create RAID or JBOD ranks
- Add volumes and reassign volumes between different hosts
- View communication resource settings, such as TCP/IP configuration and users
- View cluster LIC levels
- Select an authorization level for each user

3.1 Changes on the ESS Specialist

ESS Specialist has been modified and improved (from its previous releases) to include new functions that simplify configurations of LUNs and provide support for fibre channel adapters. In this section, we will only discuss those topics related to Fibre Channel attachment.

3.1.1 LUN masking

Assigning ESS LUNs in FC is a lot different than in SCSI. In SCSI, LUNs are assigned based on SCSI ports, independent of which hosts may be attached to those ports. So if you have multiple hosts attached to a single SCSI port (ESS supports up to four hosts per port), all of them will have access to the same LUNs available on that port.

In FC, LUN affinity is based on the world-wide port name (WWPN) of the adapter on the host, independent of to which ESS FC port the host is attached. In a switched fabric with multiple connections to the ESS, this concept of LUN affinity enables the host to see the same LUNs on different paths. If the host is not capable of recognizing that the set of LUNs seen via each path is the same, this may present data integrity problems when the LUNs are used by the operating system. To get around this problem, you can do switch zoning, or you can install the IBM Subsystem Device Driver (SDD) which is preferred. Aside from preventing the above problem, SDD also provides multipathing and load balancing which improves performance and path availability. For more information on SDD, refer to Chapter 4, “Subsystem Device Driver” on page 31.

3.1.1.1 Configuring LUNs on FC

The steps involved in configuring LUNs for Fibre Channel attachment are:

1. Define Fibre Channel host.
2. Assign volumes to the Fibre Channel host.

Unlike in SCSI, where you only need to define one host then attach ESS SCSI ports onto it, Fibre Channel requires a host icon for every HBA installed into your host, even if the adapters are installed on the same host. This is because each adapter has a unique WWPN, and hosts are defined based on this HBA’s identification.

Defining Fibre Channel host

1. Open your browser. On the location field, type in the IP address or hostname of any cluster on the ESS. This is with the assumption that both clusters have the ESS Specialist enabled (which is the default). This brings up the ESS Specialist welcome screen as seen in Figure 9.

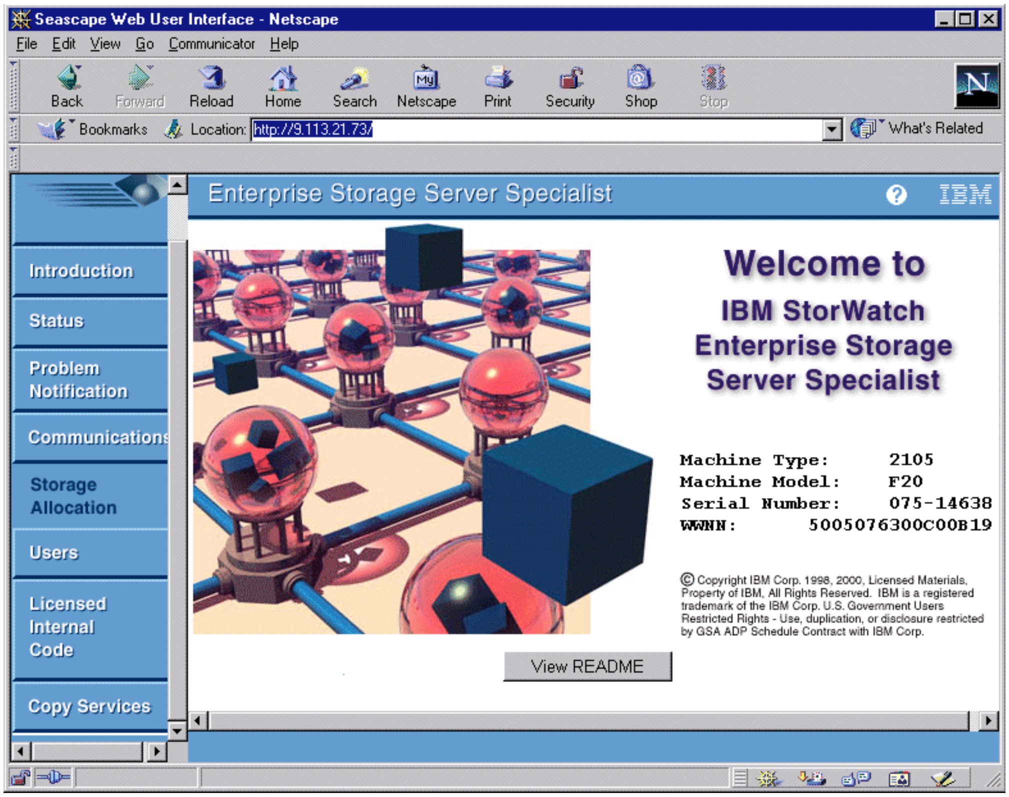

Figure 9. ESS Specialist welcome screen

2. At the left side of the screen, click **Storage Allocation**.
3. On the New Site Certificate window, click **Next** (four times), **Finish** (once) and **Continue** (once).
4. On the logon screen, type in the username and password of an administrator account. The default user is ***STORWATCH*** and the default password is ***SPECIALIST***. After you log on, you should see a similar screen as shown on Figure 10.

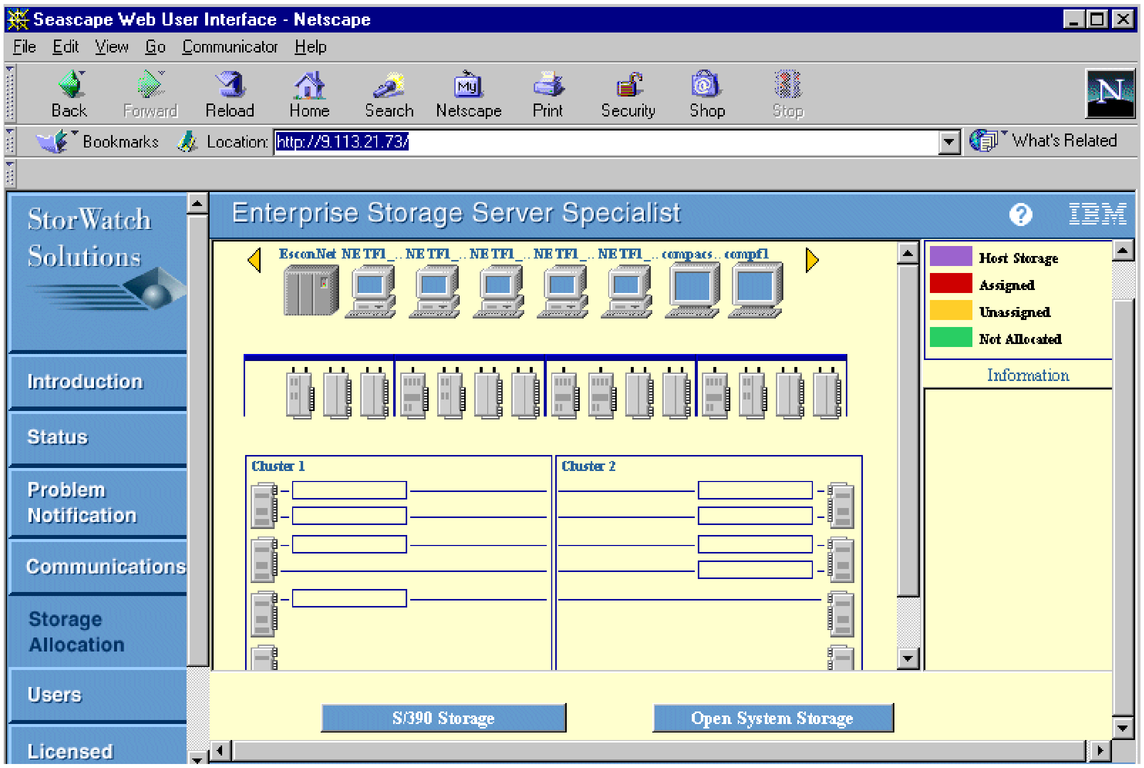

Figure 10. Storage Allocation screen

5. Click **Open System Storage** (Figure 11).

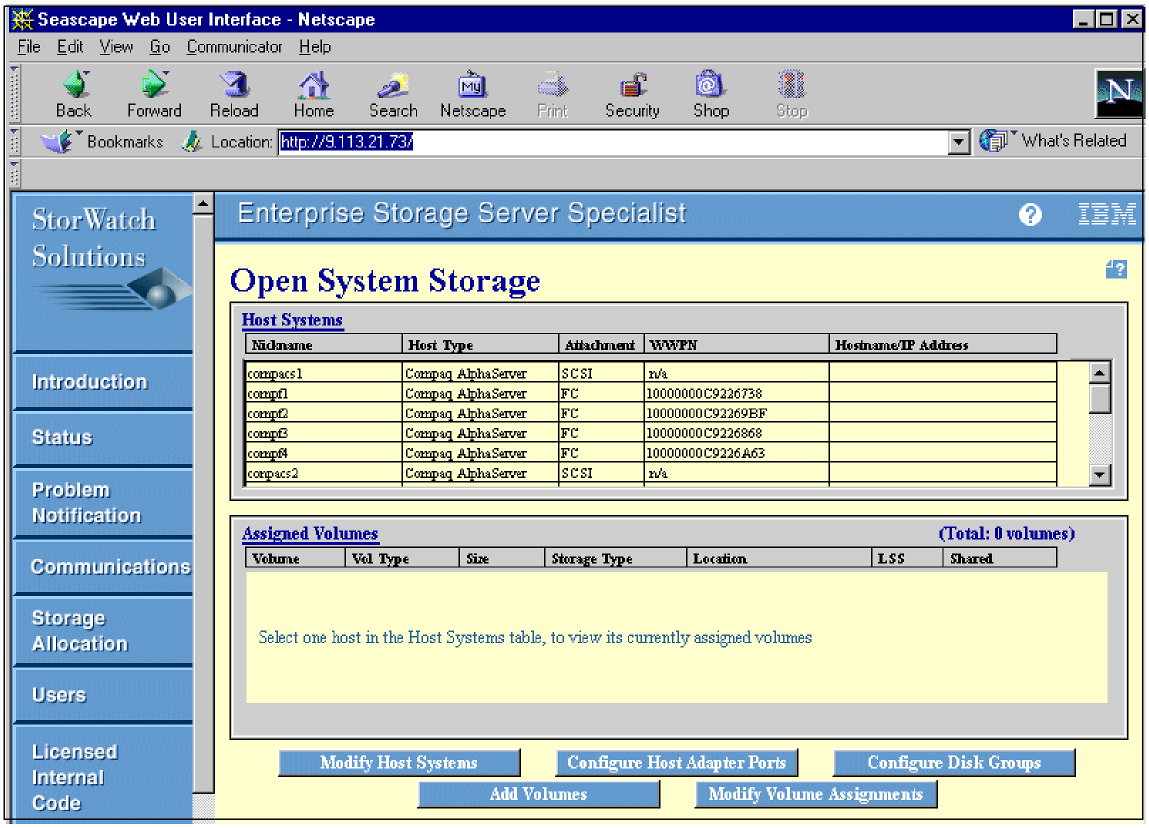

Figure 11. Open System Storage screen

6. Click **Modify Host Systems** (Figure 12).
7. Under Host Attributes, type the host's nickname on the Nickname field.
8. Select the host type using the pull-down arrow on the Host Type field.
9. Under Host Attachment field, click the pull-down arrow and select **Fibre Channel Attached**.
10. Under World-Wide Port-Name (WWPN), type in or select from the pull-down arrow the WWPN of the host bus adapter.
11. Under Hostname/IP address field, type in the name or IP of the host.
12. After completing all the fields, select **Add>>**. The nickname of the host that you just created should now appear on the Host Systems List.

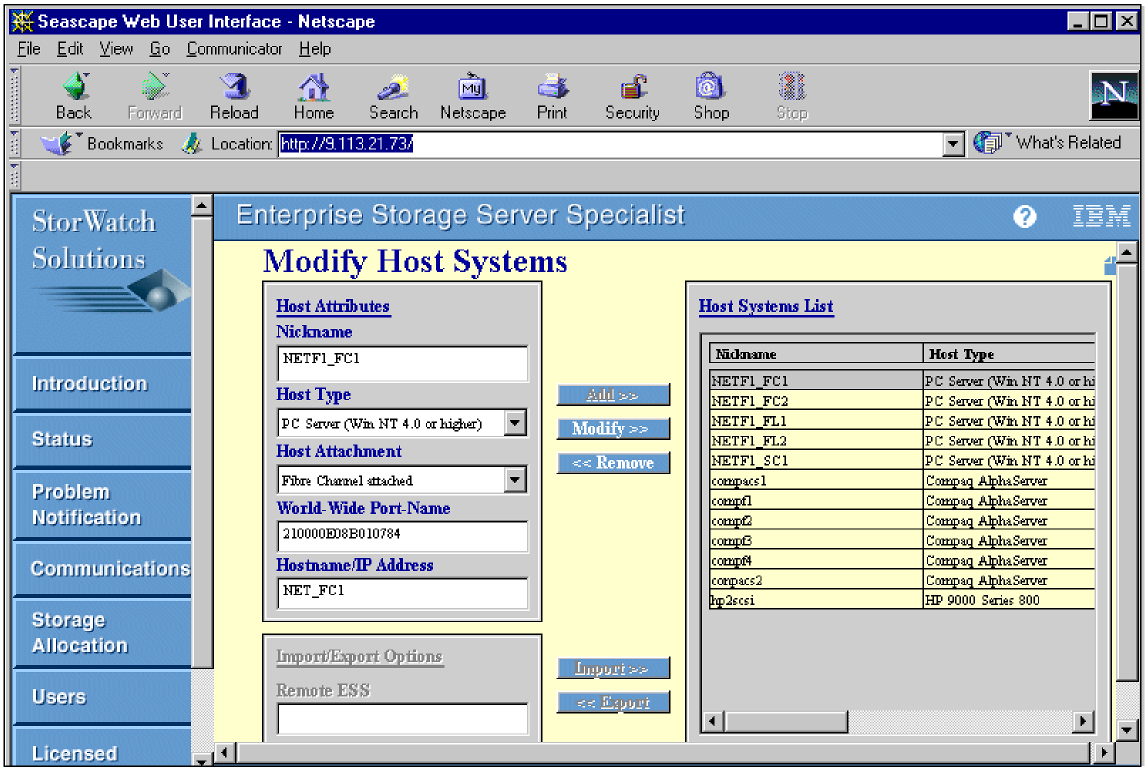

Figure 12. Modify Host Systems screen

13.Click **Perform Configuration Update** for the changes to take effect.

Now that you have defined the host, you're ready to assign volumes to it.

Assigning Volumes to Host

1. From the Main Menu, go to **Storage Allocation** (refer to Figure 9 on page 21).
2. Click **Open System Storage** (refer to Figure 10 on page 22).
3. Click **Modify Volume Assignments** (Figure 13).

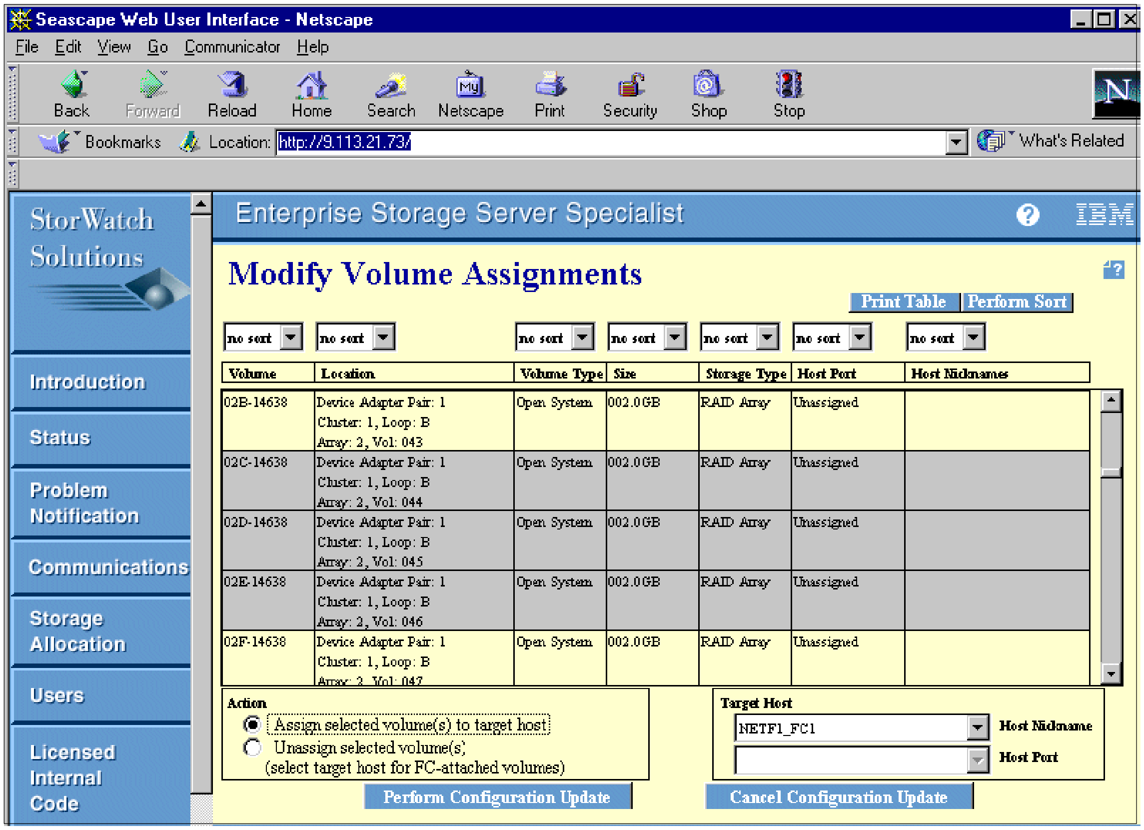

Figure 13. Modify Volume Assignments screen

4. Highlight the volumes you want to assign to the host. If multiple contiguous volumes are to be selected, highlight the first, press Shift key, then highlight the last volume. If non-contiguous, use Ctrl key, and highlight each volume. Once done, click **Perform Configuration Update**.
5. Once the progress bar ends, click **OK**. The volumes you assigned are now available to the host. However, you may need to reboot the host, or rescan the bus so that the volumes can be used by the operating system.

3.1.2 Topology changes for an ESS FC port

If your host is directly attached to the ESS, you need to make sure that they are using the same topology (either loop or point to point). Some HBAs can be configured to change the topology, some can not. On the other hand, ESS gives you the flexibility to set it according to your requirement.

3.1.2.1 Steps in changing the topology

In this example, we'll show how to change the topology from *point to point* to *loop*:

1. Go to **ESS Specialist** and logon using an administrator account.
2. Go to **Storage Allocation**.
3. Go to **Open System Storage**.
4. Select **Configure Host Adapter Ports**.
5. Select the Fibre Channel card which you want to set the topology by clicking on the adapter icon, or by using the pull-down arrow on the Host Adapter Port (Figure 14).

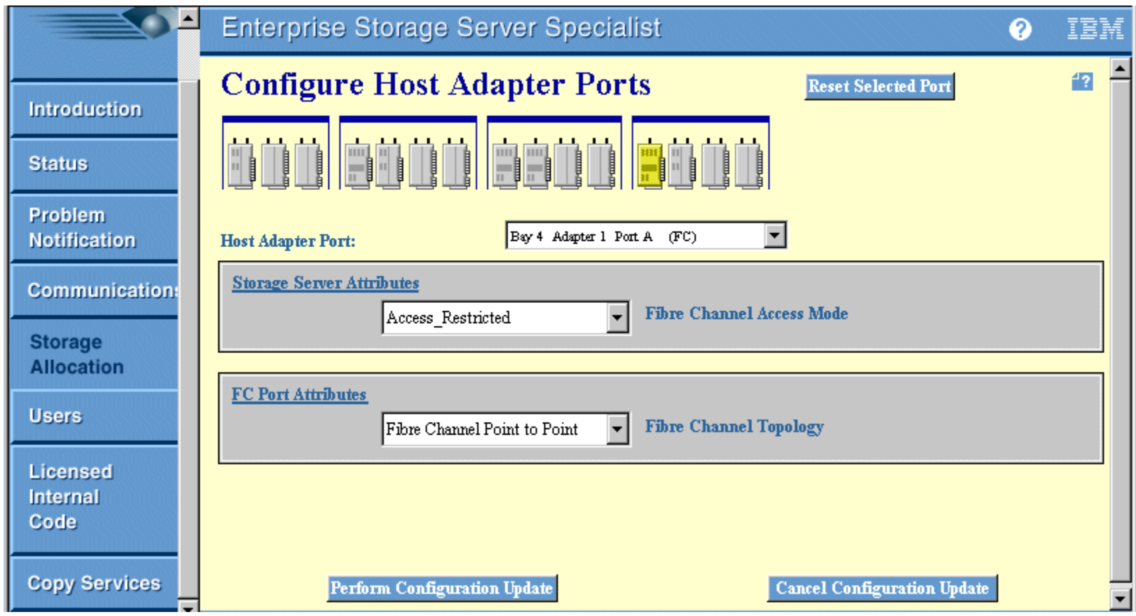

Figure 14. Window for configuring host adapter ports

6. Under FC Port Attributes, select **Undefined** (Figure 15). This puts the card in service mode, before allowing you to make any changes.

Note

Changing the FC port attributes is a disruptive action for volumes assigned to servers that have access to the ESS only through the selected port.

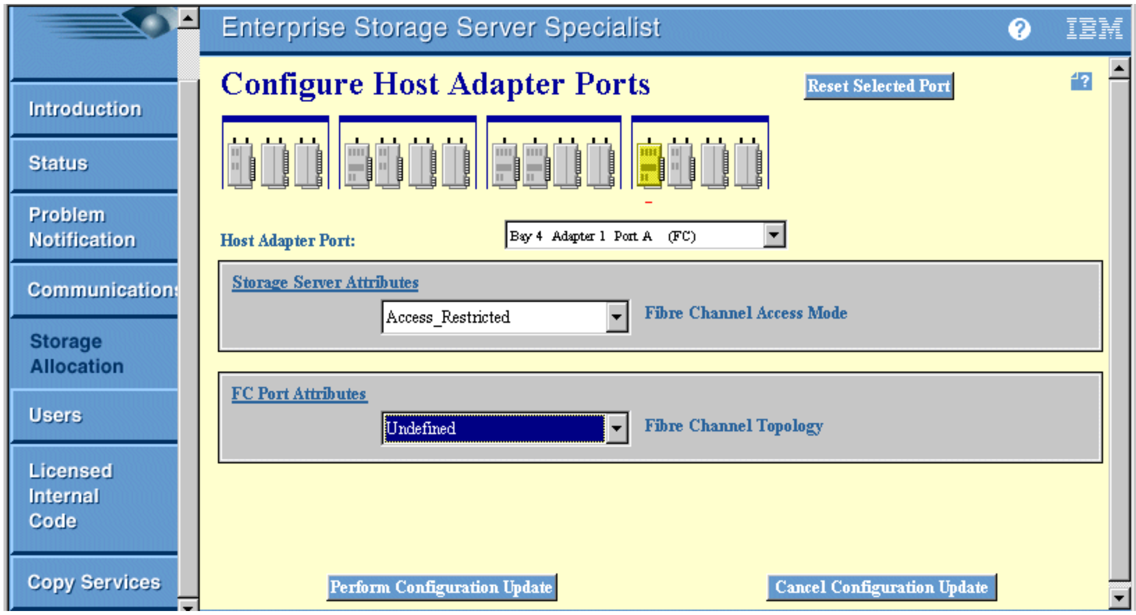

Figure 15. Putting the card in service mode

7. Click **Perform Configuration Update**. Once done, click **OK**.
8. Go again to **Configure Host Adapter Ports**.
9. Again, select the Fibre Channel card which you want to change the topology by clicking on the adapter icon, or using the pull-down arrow on the Host Adapter Port.
10. Under FC Port Attributes, select the proper topology (in this example in Figure 16, we use *Fibre Channel Arbitrated Loop).*

Note

ESS provides two choices:

1. Arbitrated Loop
2. Point-to-Point

There is no special selection for switched fabric. For a switched fabric you will select the point-to-point protocol as this is the protocol used in a fabric.

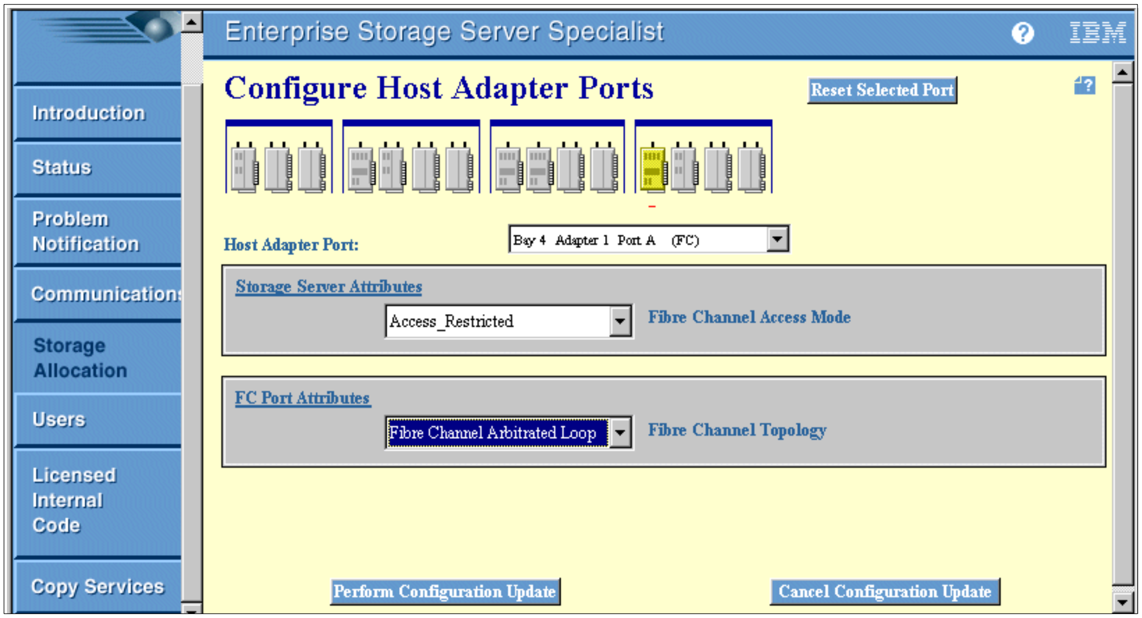

Figure 16. Selecting the Fibre Channel topology

3. Click **Perform Configuration Update**. Once done, click **OK**.

Notice that there is also the *Storage Server Attributes* field (see Figure 14 on page 26*)*. This relates to how the ESS allows host access to its volumes through all its FC interfaces. There are two kinds of Fibre Channel Access Mode. They are:

1. Access Any

   In this mode, any host's Fibre Channel adapter for which there has been no Access Profile (volumes assigned to the HBA's WWPN) defined can access all LUNs in the ESS (or the first 256 LUNs if the host does not have the "Report LUNs" capability). This would typically be used with a large server, such as a Sun UE1000 or RS6000 SP, where only one host system is attached to the storage.

- Access Restricted

  In this mode, any host's Fibre Channel adapter for which there has been no Access Profile defined *cannot* access any LUNs in the ESS.

In either mode, a host's Fibre Channel adapter with an access profile can see only those LUNs defined to it.

This attribute can only be changed by the Customer Engineer (CE) via the service port. If this is changed, the ESS needs a reboot for the change to take effect.

For more information on ESS Specialist, refer to *IBM Enterprise Storage Server Web Users Interface Guide* , SC26-7346.

Chapter 4. Subsystem Device Driver

The IBM Subsystem Device Driver (SDD) software, which resides in the host server with the disk device driver for the Enterprise Storage Server (ESS), uses redundant connections between the host server and disk storage in a ESS providing enhanced performance and data availability. These connections comprise many different components through which data flows during input and output processes. Redundancy and the ability to switch between these components provides many different paths for the data to travel. In the event of failure in one input-output path, it automatically switches to another input-output path. This automatic switching in the event of failure is called failover.

Note: You cannot run Subsystem Device Driver in an environment where more than one host is attached to the same logical unit (LUN) on a Enterprise Storage Server, for example, a multi-host environment. This restriction includes clustered hosts, such as RS/6000 servers running HACMP.

The SDD provides these functions:

- Enhanced data availability
- Automatic path failover and recovery to an alternate path
- Dynamic load balance of multiple paths
- Command line that displays device, path, and adapter information

In most cases, host servers are configured with multiple host adapters and SCSI connections to a Enterprise Storage Server that, in turn, provide internal component redundancy. With dual clusters and multiple host interface adapters, the ESS provides more flexibility in the number of input-output paths that are available. When there is a failure, the SDD reroutes input-output operations from the failed path to the remaining paths. This prevents a bus adapter on the host server, external SCSI or FC cable, or cluster or host interface adapter on the ESS from disrupting data access.

Multi-path load balancing of data flow prevents a single path from becoming overloaded, causing input-output (I/O) congestion that occurs when many I/O operations are directed to common devices along the same input-output path. Normally, selection is performed on a global rotating basis; however, the same path is used when two sequential write operations are detected.

SDD supports concurrent use of SCSI and FC paths to the same volumes on the ESS.

4.1 Configuring Subsystem Device Driver for an AIX host system

Note: You must have root access to configure the SDD and have SDD already installed.

You can configure the IBM Subsystem Device Driver on an AIX host system in one of two ways. Enter **shutdown -rF** to start your AIX host server again. This will take a few minutes. Or, you can use SMIT.

Before you configure the SDD, ensure the Enterprise Storage Server is on and configured correctly on the AIX host system, and that the dpo.ibmssd.rte software is installed on the AIX host system. Vary off (deactivate) all active volume groups with ESS subsystem disks using the varyoffvg LVM command.

Note: Before you vary off a volume group, unmount all file systems of that volume group that are mounted.

To configure the SDD, follow these steps:

1. Enter **SMIT** from your desktop window. The System Management Interface Tool Menu displays.
2. Highlight **Devices** and press Enter. The Devices menu displays.
3. Highlight **Data Path Device** and press Enter. The Data Path Device screen displays.
4. Highlight **Define and Configure All Data Path Optimizer Devices** and press Enter. Configuration begins.

To return the volume group of ESS subsystem disks to service that you varied off prior to configuring the Subsystem Device Devices, enter the **varyonvg** command. Then, mount the volume group file systems that you previously unmounted.

4.2 Verify the Subsystem Device Driver configuration on AIX host

To verify configuration of the Subsystem Device Driver on an AIX host system, follow these steps:

1. Enter **SMIT** from your desktop window. The System Management Interface Tool Menu displays.
2. Highlight **Devices** and press Enter. The Devices menu displays.
3. Highlight **Data Path Device** and press Enter. The Data Path Device screen displays.

4. Highlight **Display Data Path Device Configuration** and press Enter. A list is displayed of the condition (either Defined or Available) of all Subsystem Device Driver pseudo devices, in addition to the multiple paths of each device. If any device is listed as Defined, the configuration was not successful. Check the configuration procedure again.

To verify multiple attached paths of each adapter connected to a ESS port, follow these steps:

1. Enter **SMIT** from your desktop window. The System Management Interface Tool menu displays.
2. Highlight **Devices** and press Enter. The Devices menu displays.
3. Highlight **Data Path Device** and press Enter. The screen displays.
4. Highlight **Display Data Path Device Adapter Status** and press Enter. The screen displays all paths of each adapter as shown in Figure 17.

```
                                COMMAND STATUS

Command: OK             stdout: yes           stderr: no

Before command completion, additional instructions may appear below.

vpath90 (Avail ) 200FCA49 = hdisk91 (Avail ) hdisk101 (Avail )
vpath91 (Avail ) 201FCA49 = hdisk92 (Avail ) hdisk102 (Avail )
vpath92 (Avail ) 202FCA49 = hdisk93 (Avail ) hdisk103 (Avail )
vpath93 (Avail ) 203FCA49 = hdisk94 (Avail ) hdisk104 (Avail )
vpath94 (Avail ) 204FCA49 = hdisk95 (Avail ) hdisk105 (Avail )









F1=Help             F2=Refresh          F3=Cancel           Esc+6=Command
Esc+8=Image         Esc+9=Shell         Esc+0=Exit          /=Find
n=Find Next
```

Figure 17. SMIT output of data path device configuration

4.3 Verifying Subsystem Device Driver configuration on Windows NT

1. Log on using an administrator account.
2. Go to Disk Administrator.
3. Verify that the number of ONLINE disks equals the number of actual volumes attached to the host.
4. Verify that the number of OFFLINE disks equals the number of paths from the host to the ESS, multiplied by the number of volumes attached to the host.

4.4 SDD commands

SDD provides several commands that can be used to display information about adapters and devices (logical volumes on the ESS), and allow you to set adapters to online or offline state, which is useful for maintenance purposes. Below are some examples of these commands are shown in Figure 18.

- Datapath query adapter

  This command displays information about the adapter (SCSI port, paths).

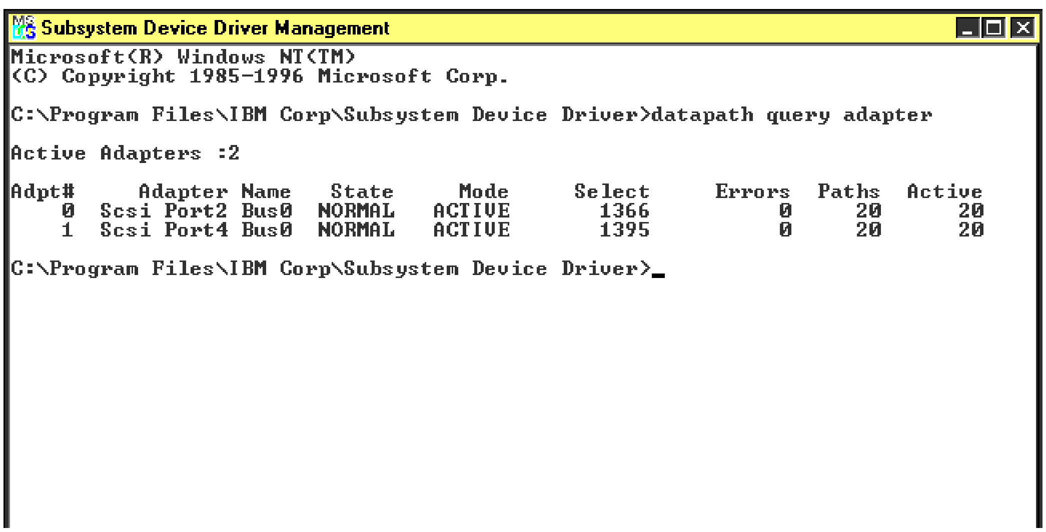

Figure 18. Datapath query adapter command output

- Datapath query device

  This command shown in Figure 19 displays information about a volume.

```
Subsystem Device Driver Management

Total Devices : 12

DEV#:   0  DEVICE NAME: Disk5 Part1  TYPE: 2105F20   SERIAL: 20214638
=====================================================================
Path#              Adapter/Hard Disk     State     Mode       Select      Errors
    0     Scsi Port2 Bus0/Disk5 Part1     OPEN    NORMAL         43           0
    1     Scsi Port4 Bus0/Disk8 Part1     OPEN    NORMAL         52           0

DEV#:   1  DEVICE NAME: Disk5 Part0  TYPE: 2105F20   SERIAL: 20214638
=====================================================================
Path#              Adapter/Hard Disk     State     Mode       Select      Errors
    0     Scsi Port2 Bus0/Disk5 Part0     OPEN    NORMAL         49           0
    1     Scsi Port4 Bus0/Disk8 Part0     OPEN    NORMAL         41           0

DEV#:   2  DEVICE NAME: Disk4 Part1  TYPE: 2105F20   SERIAL: 20114638
=====================================================================
Path#              Adapter/Hard Disk     State     Mode       Select      Errors
    0     Scsi Port2 Bus0/Disk4 Part1     OPEN    NORMAL         38           0
    1     Scsi Port4 Bus0/Disk7 Part1     OPEN    NORMAL         57           0

DEV#:   3  DEVICE NAME: Disk4 Part0  TYPE: 2105F20   SERIAL: 20114638
=====================================================================
Path#              Adapter/Hard Disk     State     Mode       Select      Errors
    0     Scsi Port2 Bus0/Disk4 Part0     OPEN    NORMAL         32           0
    1     Scsi Port4 Bus0/Disk7 Part0     OPEN    NORMAL         46           0

DEV#:   4  DEVICE NAME: Disk3 Part1  TYPE: 2105F20   SERIAL: 20014638
=====================================================================
Path#              Adapter/Hard Disk     State     Mode       Select      Errors
    0     Scsi Port2 Bus0/Disk3 Part1     OPEN    NORMAL         47           0
    1     Scsi Port4 Bus0/Disk6 Part1     OPEN    NORMAL         48           0

DEV#:   5  DEVICE NAME: Disk3 Part0  TYPE: 2105F20   SERIAL: 20014638
=====================================================================
Path#              Adapter/Hard Disk     State     Mode       Select      Errors
    0     Scsi Port2 Bus0/Disk3 Part0     OPEN    NORMAL        137           0
    1     Scsi Port4 Bus0/Disk6 Part0     OPEN    NORMAL        119           0

DEV#:   6  DEVICE NAME: Disk2 Part1  TYPE: 2105F20   SERIAL: 01614638
=====================================================================
-- More  --
```

Figure 19. Datapath query device command output

- Datapath set adapter *(online or offline). See* Figure 20 and Figure 21

  This puts the adapter online or offline, and it is useful for maintenance purposes (replacing components on a failed link).

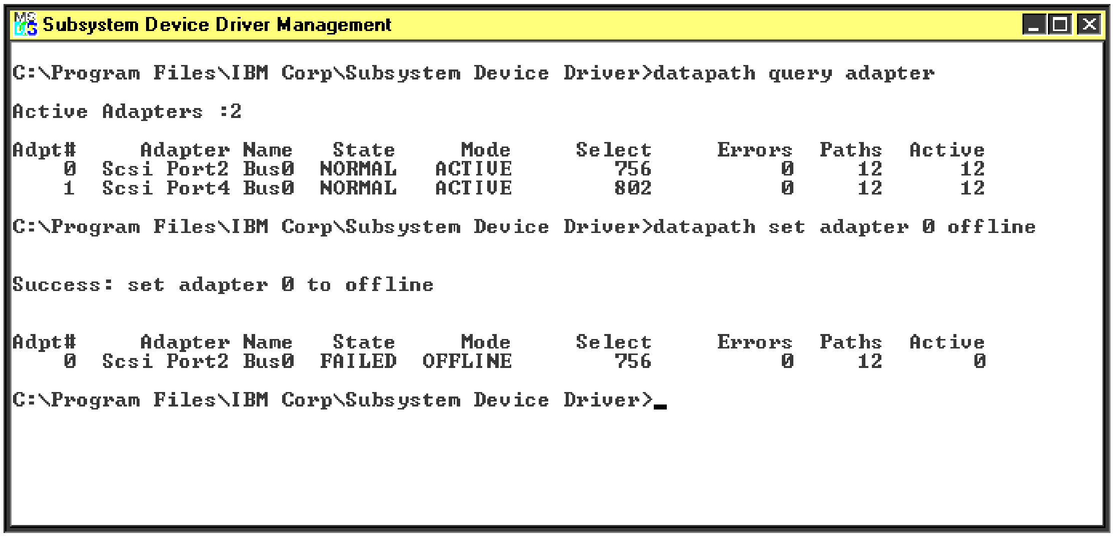

```
Subsystem Device Driver Management

C:\Program Files\IBM Corp\Subsystem Device Driver>datapath query adapter

Active Adapters :2

Adpt#     Adapter Name   State    Mode      Select     Errors  Paths  Active
    0   Scsi Port2 Bus0  NORMAL   ACTIVE       756          0     12      12
    1   Scsi Port4 Bus0  NORMAL   ACTIVE       802          0     12      12

C:\Program Files\IBM Corp\Subsystem Device Driver>datapath set adapter 0 offline


Success: set adapter 0 to offline


Adpt#     Adapter Name   State    Mode      Select     Errors  Paths  Active
    0   Scsi Port2 Bus0  FAILED  OFFLINE       756          0     12       0

C:\Program Files\IBM Corp\Subsystem Device Driver>_
```

Figure 20. Putting the adapter offline

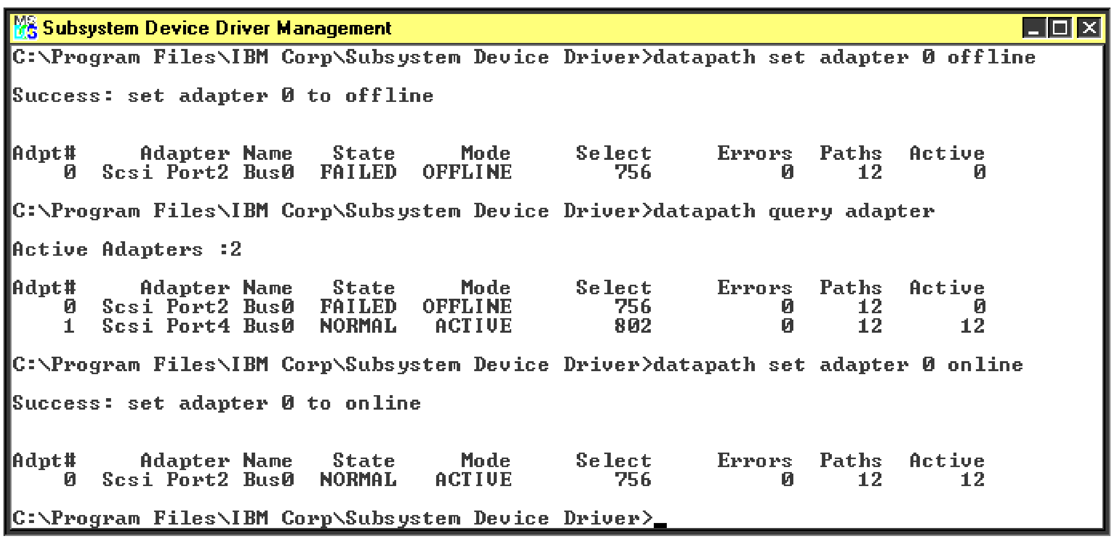

```
Subsystem Device Driver Management
C:\Program Files\IBM Corp\Subsystem Device Driver>datapath set adapter 0 offline

Success: set adapter 0 to offline


Adpt#     Adapter Name   State    Mode      Select     Errors  Paths  Active
    0   Scsi Port2 Bus0  FAILED  OFFLINE       756          0     12       0

C:\Program Files\IBM Corp\Subsystem Device Driver>datapath query adapter

Active Adapters :2

Adpt#     Adapter Name   State    Mode      Select     Errors  Paths  Active
    0   Scsi Port2 Bus0  FAILED  OFFLINE       756          0     12       0
    1   Scsi Port4 Bus0  NORMAL   ACTIVE       802          0     12      12

C:\Program Files\IBM Corp\Subsystem Device Driver>datapath set adapter 0 online

Success: set adapter 0 to online


Adpt#     Adapter Name   State    Mode      Select     Errors  Paths  Active
    0   Scsi Port2 Bus0  NORMAL   ACTIVE       756          0     12      12

C:\Program Files\IBM Corp\Subsystem Device Driver>_
```

Figure 21. Putting the adapter back online

For more information on SDD, refer to the html file (dpo1mst.htm) that comes with the SDD installer.

Chapter 5. Configuring host systems with Fibre Channel adapter

This chapter describes how to configure various hosts with the native Fibre Channel adapter directly attached to the Enterprise Storage Server (ESS).

5.1 How to find World Wide Port Name (WWPN)

This section describes what WWPN is, how to determine the WWPN on various hosts, and how it is used on ESS.

The WWPN is a burnt-in address on the Fibre Channel adapter card the same as the MAC address on a network adapter card. It is a 16-digit long address and used to assign LUNs on ESS for the host. Once the Fibre Channel adapter is installed on a host, you need to identify WWPN; it is very important to know how to figure out the WWPN. The following sections describe how to find the WWPN on various host systems.

5.1.1 On an AIX host

You need to enter `lscfg | grep fcs` at the AIX command prompt to display how many Fibre Channel adapters are installed on your host. Figure 22 shows how the display looks.

```
dusenberg> lscfg | grep fcs
+ fcs1            20-58            FC Adapter
+ fcs3            20-60            FC Adapter
+ fcs0            10-68            FC Adapter
dusenberg> █
```

Figure 22. Installed Fibre Channel adapters

To find the WWPN address for the Fibre Channel adapter installed on your system, you need to enter `lscfg -vl fcs0`, where `fcs0` is an adapter number. It will display a detailed output with a network address which is your WWPN address for that Fibre Channel adapter. You will need to do the same for each adapter to get the WWPN address. Figure 23 shows details for *adapter fcs0*.

```
dusenberg> lscfg -vl fcs0
  DEVICE            LOCATION          DESCRIPTION

  fcs0              10-68             FC Adapter

        Network Address.............10000000C9218C9C
        ROS Level and ID............02903215
        Device Specific.(Z0)........4002206D
        Device Specific.(Z1)........10020193
        Device Specific.(Z2)........3001506D
        Device Specific.(Z3)........02000909
        Device Specific.(Z4)........FF101370
        Device Specific.(Z5)........02903215
        Device Specific.(Z6)........06113215
        Device Specific.(Z7)........07113215
        Device Specific.(Z8)........10000000C9218C9C
        Device Specific.(Z9)........SS3.20A5
        Device Specific.(ZA)........S1F3.20A5
        Device Specific.(ZB)........S2F3.20A5
        Device Specific.(YL)........P2-I3

dusenberg>
```

Figure 23. Detailed output of a Fibre Channel adapter

5.1.2 On a Windows NT Server

To find the WWPN for Qlogic Card, QLA2100F & QLA2200F:

1. Boot the server.
2. Once the Qlogic BIOS appears, press **Alt + Q** to go to Fast!UTIL.
3. If you have multiple adapters installed, go to **Select Host Adapter**, and select whichever HBA you want to get the WWPN. Press Enter.
4. Go to **Configuration Settings** and press Enter.
5. Select **Host Adapter Settings** and press Enter.
6. Look for the Adapter Node Name field. Beside it is the WWPN (16 characters in length) for the HBA.

   For example: `200000E08B00CC5B`

Note

For hosts with multiple adapters, you will need to identify which adapter is inserted in which slot, to be able to correctly map the corresponding WWPN. This normally depends on the scanning of the PCI bus. Qlogic BIOS Utility presents the order of the HBAs in how the bus is scanned. Consult your PC server hardware documentation for further related information.

Another method is to add the cards one at a time. After getting the WWPN for the first HBA, insert the second (no need to remove the first), boot the server, then get the WWPN, and so on until you finish all the cards. This requires much more time and effort though, but lessens the possibility of mismatched WWPN.

5.1.3 On an HP-UX host

You need to enter `ls -l /dev/fcms*` at the HP-UX command prompt to display all Fibre Channel adapters that are installed on your host. Figure 24 shows how the display looks.

```
# ls -l /dev/fcms*
crw-rw-rw-   1 bin       bin        143 0x010000 Mar  9 15:39 /dev/fcms1
# █
```

Figure 24. Fibre Channel adapter device display on HP-UX

Once you find out all Fibre Channel adapters installed on your system, you will need to enter `fcmsutil /dev/fcms1`, where `fcsm1` is an adapter number. It will display detailed output with a network address, which is your WWPN address. You will need to do the same for each adapter to get the WWPN address. Figure 25 shows how the output looks.

```
# fcmsutil /dev/fcms1

                    Local N_Port_ID is = 0x000000
          N_Port Node World Wide Name = 0x1000001083FC1157
          N_Port Port World Wide Name = 0x1000001083FC1157
                             Topology = UNKNOWN
                                Speed = 1062500000 (bps)
                          HPA of card = 0xFA6FF000
                          EIM of card = 0xFFFFFFFF
                         Driver state = OFFLINE
               Number of EDB's in use = 0
               Number of OIB's in use = 0
  Number of Active Outbound Exchanges = 0
      Number of Active Login Sessions = 0

# █
```

Figure 25. Detailed output of fcmsutil command with WWPN information

5.2 Configuring Fibre Channel on an AIX host

In this section we discuss how to configure disks attached via the native Fibre Channel adapter. Before you attach fibre cable to the ESS unit you need to make sure your AIX server has proper level of PTFs, device driver, microcode and 2105 SDD host attachment drivers to work the disk configuration properly.You also need to make sure that ESS configuration sheets are completed, which will be implemented on ESS to assign LUNs to your host.

To configure disks, you can use *one* of the following three procedures:

1. Type `cfgmgr` at the command prompt and press enter as root user. This will check all devices are connected after the last reboot.

   OR

2. Enter `SMIT` from your desktop window. The System Management Interface Tool Menu displays.

   a. Highlight **Devices** and press Enter. The Devices menu displays.
   b. Highlight **Install/Configure Devices Added After IPL** and press Enter.
   c. Press Enter with defaults.

OR

3. Reboot your server.

5.3 Verify configuration of the disks and availability

To verify configuration of the disks on an AIX host system, follow the steps below.

1. Enter `SMIT` from your desktop window. The System Management Interface Tool Menu displays.
2. Highlight **Devices** and press Enter. The Devices menu displays.
3. Highlight **Fixed Disk** and press Enter. The Fixed Disk screen displays.
4. Highlight **List all defined disks** and press Enter. A list is displayed of the condition (either Defined or Available) of all disks devices. If any device is listed as Defined, the configuration was not successful. Check the configuration procedure or adapter status, cables, and so on again. The screen output will look as shown in Figure 26.

```
                                COMMAND STATUS

Command: OK            stdout: yes           stderr: no

Before command completion, additional instructions may appear below.

[MORE...91]
hdisk91  Available 10-68-01      IBM FC 2105F20
hdisk92  Available 10-68-01      IBM FC 2105F20
hdisk93  Available 10-68-01      IBM FC 2105F20
hdisk94  Available 10-68-01      IBM FC 2105F20
hdisk95  Available 10-68-01      IBM FC 2105F20
hdisk96  Available 10-68-01      IBM FC 2105F20
hdisk97  Available 10-68-01      IBM FC 2105F20
hdisk98  Available 10-68-01      IBM FC 2105F20
hdisk99  Available 10-68-01      IBM FC 2105F20
hdisk100 Available 10-68-01      IBM FC 2105F20
hdisk101 Available 20-58-01      IBM FC 2105F20
hdisk102 Available 20-58-01      IBM FC 2105F20
[MORE...104]

F1=Help             F2=Refresh          F3=Cancel           Esc+6=Command
Esc+8=Image         Esc+9=Shell         Esc+0=Exit          /=Find
n=Find Next
```

Figure 26. Screen output of SMIT

5.4 Configuring Fibre Channel on a Windows NT host

In an NT environment, there are some configurations required for the system to have optimum performance. The exact configurations per environment depend on many variables, from application to the hardware, and may differ from system to system, and require a lot of knowledge on the end-to-end components. Some recommendations that help in achieving optimum system performance are shown in Table 2.

5.4.1 Configuring the Qlogic Adapter

1. Boot the server.
2. Once the Qlogic BIOS appears, press **Alt + Q** to go to Fast!UTIL.
3. If you have multiple adapters installed, go to **Select Host Adapter**, and select whichever HBA you want to configure. Press Enter.
4. Go to **Configuration Settings** and press Enter.
5. Select **Host Adapter Settings** and press Enter.
6. Compare the values and change to match the values recommended:

Table 2. Recommended values for host adapter settings

Parameter	Default	Recommended	Remarks
Host Adapter BIOS	Disabled	Disabled	This should *only* be enabled if you're booting from a disk attached to the fibre HBA
Frame Size	1024	2048	
Loop Reset Delay	5	5 (minimum)	
Adapter Hard Loop ID	Disabled	Disabled	
Hard Loop ID	0	Disabled	

7. Go back to **Configuration Settings** and select **Advanced Host Adapter Settings**.
8. Compare the values and change to match the recommended values as shown in Table 3.

Table 3. Recommended values for advanced host adapter settings

Parameter	Default	Recommended	Remarks
Execution Throttle	16	240	
Fast Command Posting	Enabled	Enabled	
>4GB Addressing	Disabled	Disabled	Should be disabled for 32 bit systems
LUNS per target	8	0	
Enable LIP reset	No	No	
Enable LIP Full Login	Yes	No	
Enable Target Reset	No	Yes	
Login Retry Count	8	20 (minimum)	
Port Down Retry Count	8	20 (minimum)	
Drivers Load RISC Code	Enabled	Enabled	
Enable Database Updates	No	No	
Disable Database Load	No	No	applicable only to QLA2200F
IOCB Allocation	256	256	
Extended Error Logging	Disabled	Disabled	may be enabled for debugging purposes

9. Press Esc twice and Save changes. Then Exit and reboot.

5.4.2 Configuring Windows NT 4.0

Windows NT 4.0 supports eight LUNs per target, achieving a maximum of 120 devices per adapter. However, most Fibre Channel HBA vendors use some kind of abstraction from the host adapter (could be in firmware but most often in the device driver). A common method is to create additional “pseudo” target IDs. This enables Windows NT to accommodate an additional 15 targets (that could have eight LUNs each) for a maximum of 240 devices.

5.4.2.1 Editing the NT Registry to support more than eight LUNs

Early Qlogic HBA device drivers had support for large LUNs disabled. This is defined in their OEMSETUP.INF, and is written to the Registry upon installation. So if you have more than eight LUNs defined, those in excess of eight will not be seen by the host. Newer device drivers had this enabled, so you won't normally need to edit it. It's good to check it however, just to make sure that it is enabled.

Here is the procedure in editing the Registry to support more than eight LUNs (Figure 27):

1. Go to **Start**, **Run**, and type `regedt32`.
2. Go to **HKEY_LOCAL_MACHINE on Local Machine**.

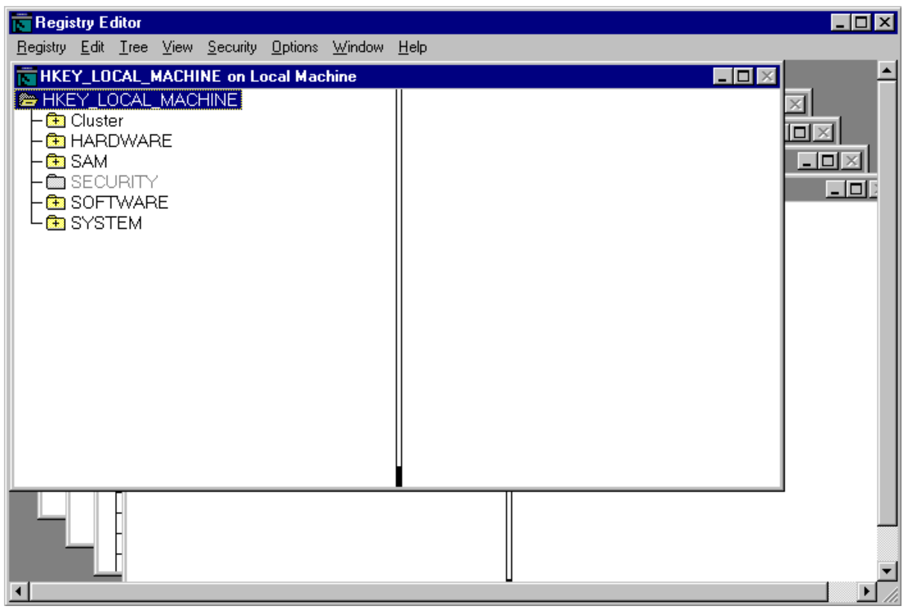

Figure 27. Registry Editor of NT

3. Double-click **System**.
4. Double-click **CurrentControlSet**.
5. Double-click **Services**.
6. Double-click **ql2100**.
7. Double-click **Parameters**.
8. Select **Devices**.

9. On the right window, look for *Large LUNs*. This should have a value of *0x1*. If it has a value of *0*, you need to edit it (Figure 28).

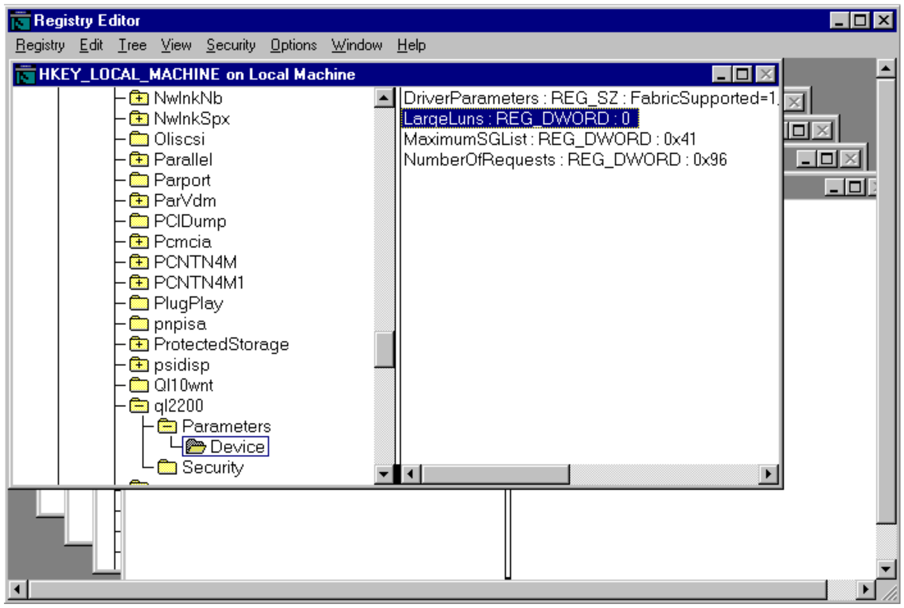

Figure 28. Location of the LargeLUNs parameter for the Qlogic adapter

10.To edit, highlight it and go to **Edit**, then **DWORD** (Figure 29**)**.

11.Under *Data*, type in **1**.

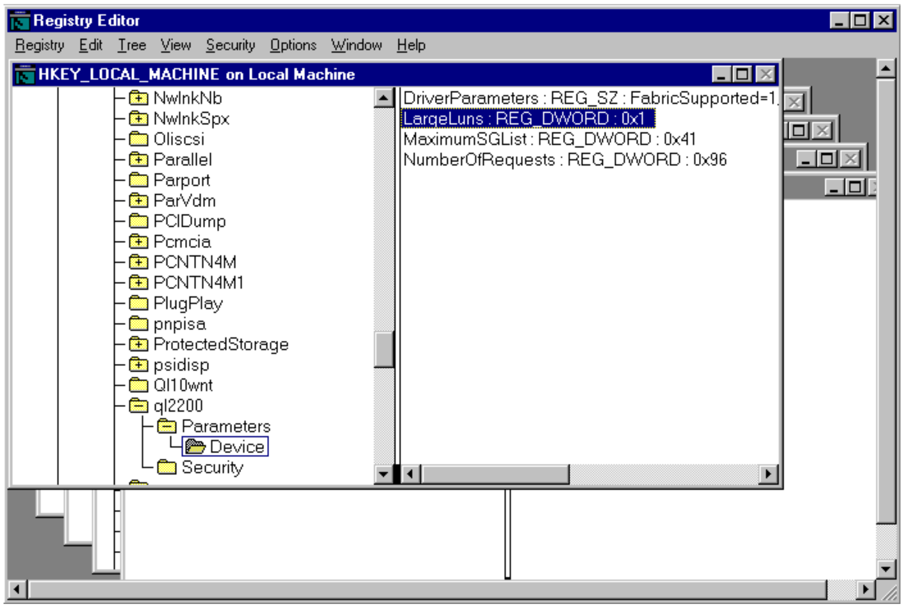

Figure 29. Editing the registry entry for LargeLUNs

12.Click **OK** and close the Registry editor.

13.Reboot for the change to take effect.

5.5 How to configure a Fibre Channel on a HP host

This section shows you how the disks look on your host once its HP host has a Fibre Channel adapter installed and connected to the ESS unit with some LUNs assigned to your host. Use the following procedure to see how all those disks look on your host.

1. Enter `sam` at your command prompt and press Enter.
2. Press enter at the next screen which displays how to use keys when using sam's menu screens.
3. Select **Disks and FIle Systems** and press Enter.
4. Select **Disk Devices** and press Enter.

The output screen will look like Figure 30.

Note

Please notice that in the following screen all new disks appear as **unused** with description of **IBM 2105F20**.

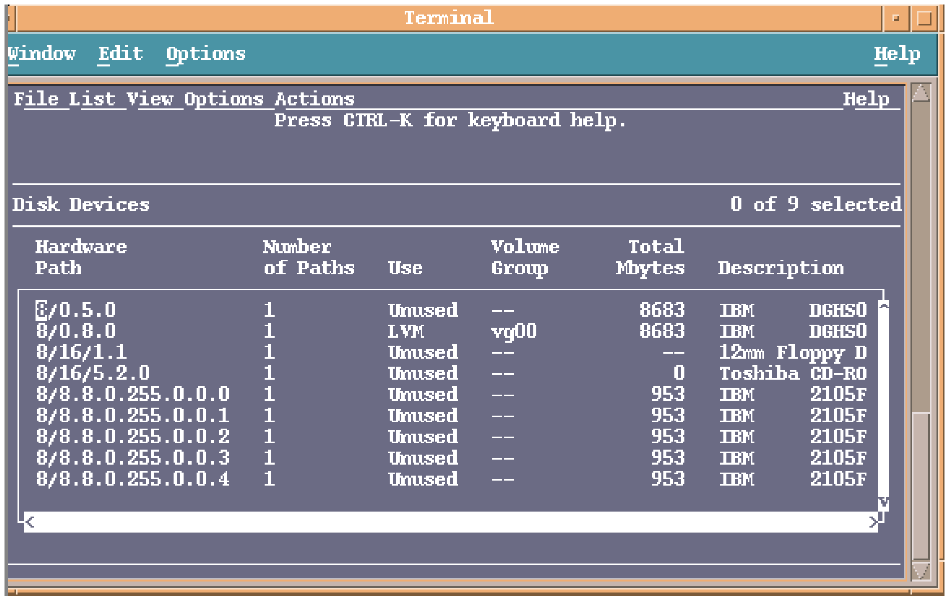

Figure 30. Screen output of disks devices

5.6 How to configure a Fibre Channel on a Sun host

To attach the ESS to a Sun host system, you must change the Sun system kernel. Make the following revisions to the sd.conf file that you find in the /kerner/drvsubdirectory to see the logical unit numbers per target ID.

1. Change to the directory by typing: cd /kernel/drv
2. Backup the sd.conf file in this directory.
3. See Figure 31 for an example of how to edit the sd.conf file. Add lines for logical unit number or target pair logical unit number or target pair. Note that the example shows how Sun host systems use logical unit number address 0 through 7 for target 8. Check the file to see if you have defined logical unit number 0 for the target ID you want to use. If logical unit

number target 0 is in use, do not repeat the logical unit number 0 definition.

```
name="sd" class="scsi"
target=8 lun=0;
name="sd" class="scsi"
target=8 lun=1;
name="sd" class="scsi"
target=8 lun=2;
name="sd" class="scsi"
target=8 lun=3;
name="sd" class="scsi"
target=8 lun=4;
name="sd" class="scsi"
target=8 lun=5;
name="sd" class="scsi"
target=8 lun=6;
name="sd" class="scsi"
target=8 lun=7;
```

Figure 31. Example of how to edit the sd.conf file

Note: On a Solaris 7 host system, you do not see the OK prompt until you type the **shutdown** command.

4. Shut down the Sun host system.
5. Start your host system by typing the following kernel re-configuration command:

   `reboot -- -r`

5.6.1 Setting the parameters for a Sun host system

Use the following parameters for a Sun host system.

- v sd_max_throttle

  This parameter specifies the maximum number of outstanding commands that the SCSI target driver (sd) attempts to queue to the targets. The default value is 256, but this parameter must be set to a value less than the maximum queue depth of any LUN. To attach to an ESS, set the value with the following formula:

  190/luns per adapter = sd_max_throttle

- v maxphys

This parameter specifies the maximum number of bytes you can be transfer for each SCSI transaction. Set this parameter in the /etc/system file as follows:

set maxphys = nnn

The default value is 126976 (124KB). The value of this parameter might vary depending on the version of the sd driver. If an I/O block size exceeds the default value, the I/O request is broken into more than one request. Each request does not exceed the default size. The value should be tuned to the intended use and application requirements.

- v sd_io_time

  This parameter sets the timeout for disk operations. The following is an example of how to set the timeout for disk operations.

  set dd:sd_io_time=0xF0

- v sd_retry_count

  This parameter sets the retry count.

  set sd:sd_retry_count=10

Chapter 6. Migrating from SCSI to Fibre Channel

In this chapter we describe the process of how to migrate disks / logical volumes within the ESS, from being SCSI attached to being Fibre Channel (FC) attached. We do this for various operating system platforms.

We assume that there is at least one ESS with SCSI attachment already installed and in use within your Data Processing Center (DP-Center).

6.1 Non-concurrent migration

This chapter describes the process of migrating disks using a non-concurrent method. This method is non-concurrent from the operating system point of view. Other servers attached to the same ESS can continue their work while this procedure is performed.

The migration procedures described in the following sections are the one we have successfully tested. There are various other ways disks can be migrated to native Fibre Channel from SCSI.

The basic assumptions used here are:

- Your server has disks that are located on an ESS in use. They are attached to the server with one SCSI interface.
- The server has one or more FC adapters, so called Host Bus Adapters (HBA), installed, the ESS has FC adapter (Feature Code 3022) installed, and the interconnections between them has been established.
- All appropriate software prerequisites, drivers and PTFs have been implemented.
- The following procedure will be executed by an experienced system administrator for their area of expertise related to the operating system. There are some commands provided in the procedure which can be used.

6.1.1 Migration from SCSI to FC general overview

Non-concurrent migration means that your ESS has disks/volumes assigned to a sever using a SCSI interface. Figure 32 shows an example.

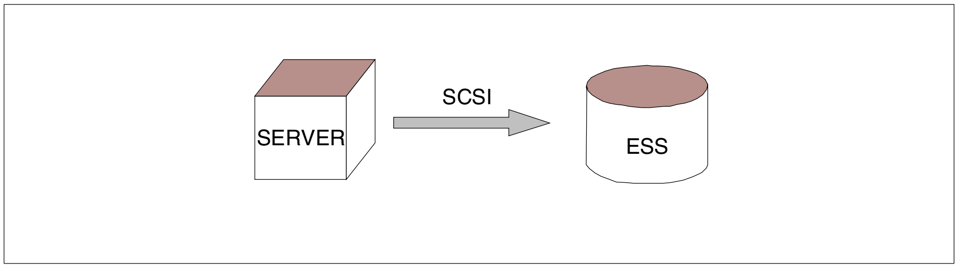

Figure 32. SCSI connection — server to ESS

Here are the steps involved in the migration process:

1. Stop the applications that are using the disks.
2. Perform all operating system required actions to “free” the disks. This could be for example:
   - Unmount file systems.
   - Vary off volume group.
   - Export volume group.
3. Unassign the disks using the ESS Specialist (Figure 33).

Note

This makes the volumes inaccessible from the server. Of course they are still available inside the ESS, and the data on these volumes is kept, too.

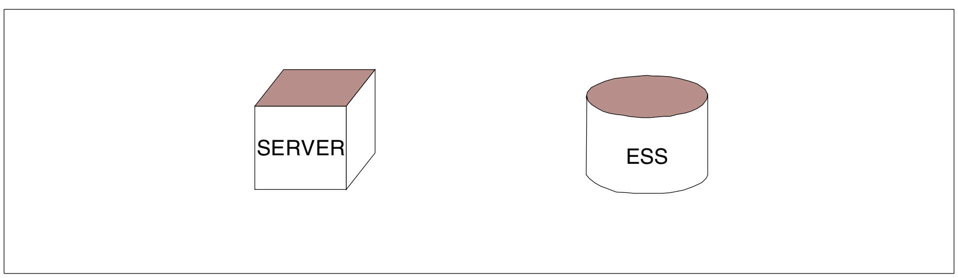

Figure 33. Unassigned disk

4. The ESS Specialist is now used to assign these disks to Fibre Channel (Figure 34).

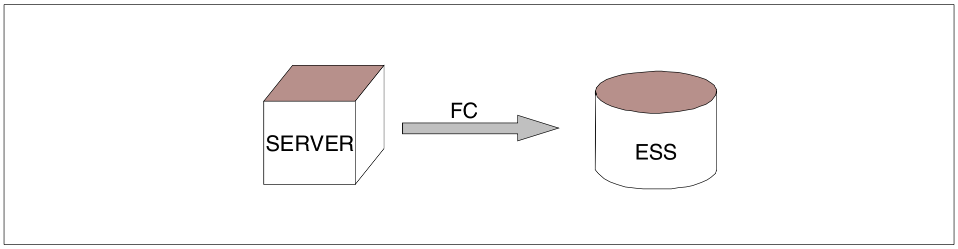

Figure 34. Fibre Channel connection — server to ESS

5. Now perform the operating system required actions to make the volumes usable by the operating system. That can involve:
   - Reboot or on AIX run Configuration Manager using command line entry (cfgmgr) or SMIT.
   - Import volume group.
   - Mount file system.

6.1.2 Migrating SCSI to FC on AIX server

This section describes the steps required on AIX to migrate disks/volumes on ESS configured from SCSI attached disks to the Fibre Channel attached disks. We have successfully tested and performed the following procedures in our test lab. Here is the procedure we used:

1. Shutdown databases and/or applications using the disks to be migrated.
2. Unmount file systems (umount <file system>).
3. Identify all disks assigned to the volume group affected (lsvg -p <volumegroup name>). We ran the following procedure to identify all disks and got serial numbers associated with the them to identify in the ESS unit.
   a. lsvg -p <volume group> | grep hdisk | cut -f1 -d" " > /tmp/disk1.
   b. Execute the script seen in Figure 35.

```
for i in `cat /tmp/disk1`
do
SN=`lscfg -vl $i | grep Serial`
echo $i, $SN >> /tmp/output
done
/#
```

Figure 35. Sample script to get Hdisk and serial number

c. The output of the script above looks like Figure 36.

```
hdisk101, Serial Number................017FCA49
hdisk102, Serial Number................018FCA49
hdisk103, Serial Number................019FCA49
hdisk104, Serial Number................01AFCA49
hdisk105, Serial Number................01BFCA49
hdisk106, Serial Number................01CFCA49
hdisk107, Serial Number................01DFCA49
hdisk108, Serial Number................01EFCA49
hdisk109, Serial Number................01FFCA49
hdisk110, Serial Number................734FCA49
hdisk111, Serial Number................735FCA49
hdisk112, Serial Number................736FCA49
hdisk113, Serial Number................634FCA49
hdisk114, Serial Number................635FCA49
hdisk115, Serial Number................636FCA49
hdisk116, Serial Number................534FCA49
hdisk117, Serial Number................535FCA49
hdisk118, Serial Number................536FCA49
hdisk119, Serial Number................434FCA49
hdisk120, Serial Number................435FCA49
hdisk121, Serial Number................436FCA49
hdisk122, Serial Number................334FCA49
```

Figure 36. Hdisk and SN information to identify on ESS unit

4. Varyoff volume group (varyoffvg <volume group name>).
5. Export volume group (exportvg <volume group name>).
6. Unassign all disks (LUNs) on ESS which are identified in Figure 36. This is done by using the ESS Specialist. A detailed explanation can be found in 3.1.1.1, “Configuring LUNs on FC” on page 20.

7. Delete all disks that belong to the volume group just exported (rmdev -l ‘hdisk1‘ -d).
8. Assign all those disks (documented in step 3 on page 53) on ESS to the host system that should get access to them. Use WWPN to identify the HBA. See 5.1, “How to find World Wide Port Name (WWPN)” on page 37 for more details.
9. Run **cfgmgr** at the AIX command prompt, or execute function **Install/Configure Devices added After IPL** using SMIT, or reboot the server.
10. After successful command execution or reboot, check to see all disks are available which were part of the volume group (compare to the document created by step 3 on page 53).
11. Once you see all disks are in *available* state you can import volume group (importvg <volumegroup> hdisk1).
12. You should get a command prompt in few seconds. Now you can verify that all disks are available (imported) to use in volume group (lsvg -p <volumegroup>).
13. Now you can mount all file systems which were unmounted in step 2 on page 53.
14. Start your application or databases.

6.1.3 Migrating SCSI to FC on Windows NT server

This section describes the steps required on Windows NT to migrate disks/volumes on ESS configured from SCSI attached disks to the Fibre Channel attached disks.

For this example, we use the setup shown in Figure 37.

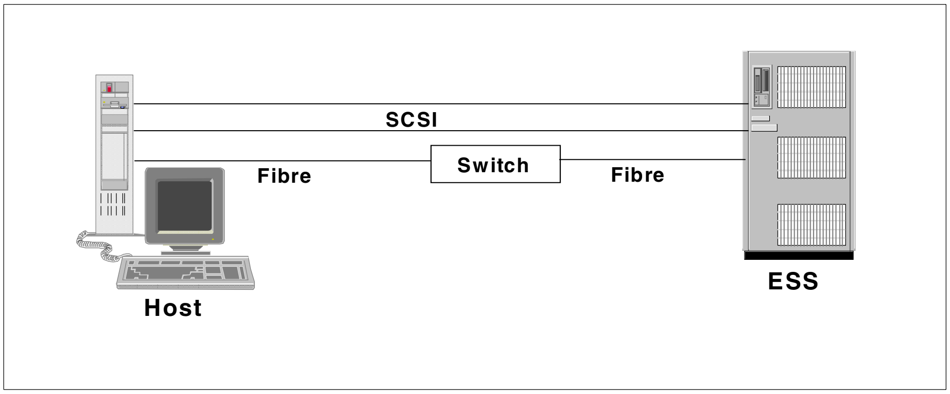

Figure 37. SCSI setup between the NT host and ESS

Software utilities used:

- Browser (Netscape or Internet Explorer)
- SAN Explorer

  SAN Explorer can be downloaded at `http://www.ibm.com/storage`
- Disk Administrator

Using SAN Explorer (or SCSI Adapter, under Control Panel) is how the disks are initially presented to the host (Figure 38).

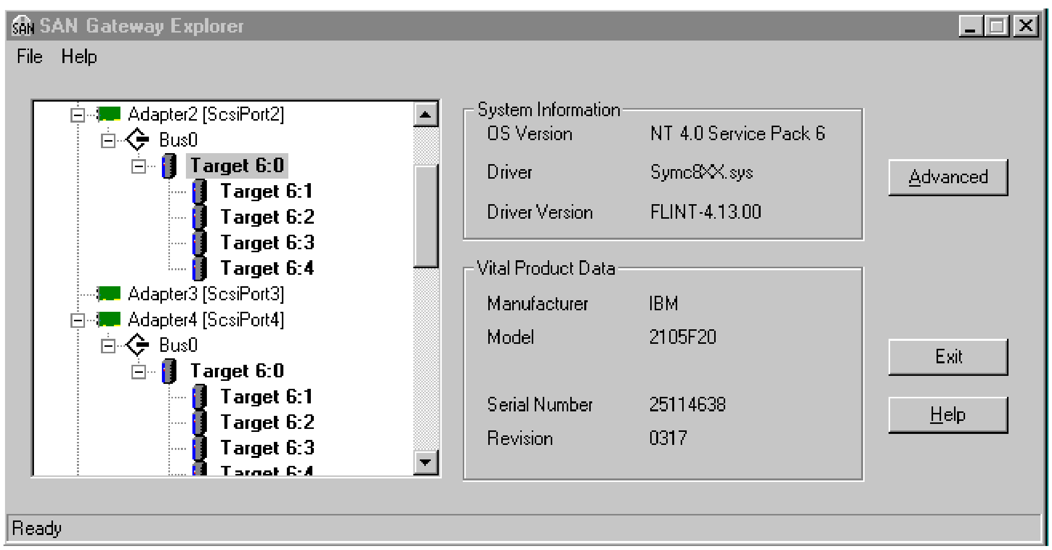

Figure 38. Initial setup of volumes attached to SCSI adapters on the host

The disks were already labeled and the drive letters have previously been assigned. This is how they would look like under Disk Administrator (Figure 39).

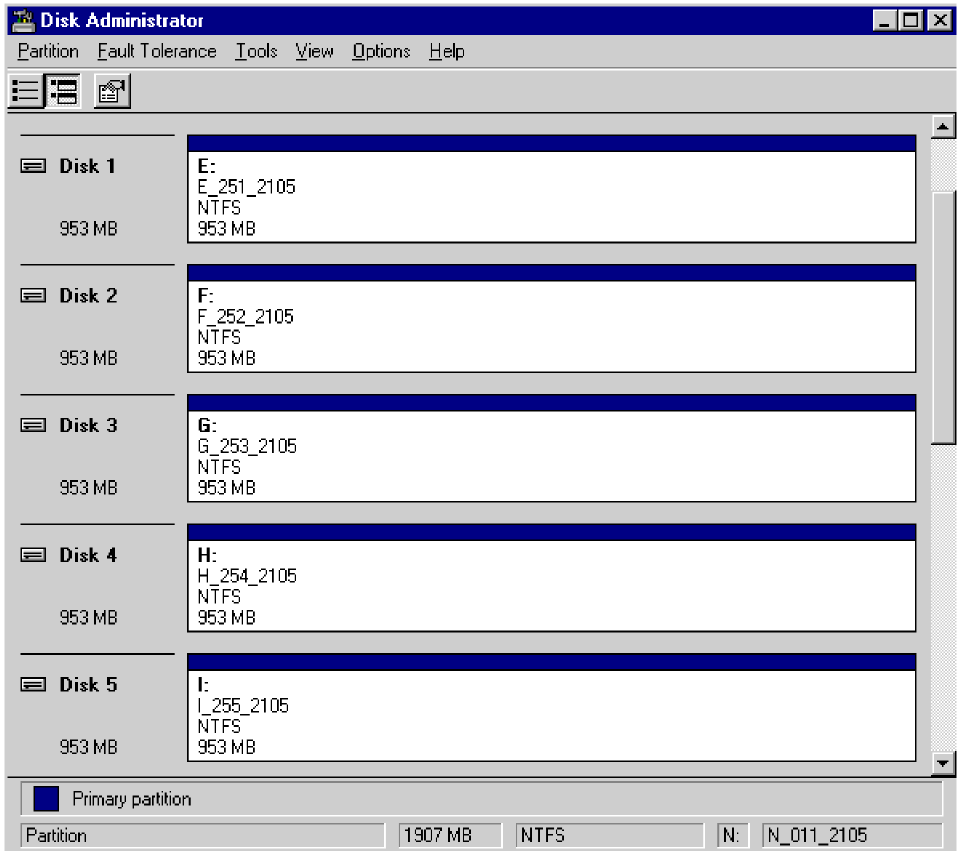

Figure 39. Disk Administrator view of the initial setup

Disk Administrator shows the volumes (only five out of ten are shown here because of the limit of Disk Administrator window). Notice that we labeled the volumes using their logical drive letter and Volume identifier (three characters in length) as shown on ESS Specialist.

Table 4 shows the relationship between the volumes and their mapping.

Table 4. Volume mapping before migration

Volume ID	SCSI Port on the ESS	SCSI Target ID & LUN	Disk Administrator Logical Drive Assignment	Disk Administrator Disk Number
251	A	6,0	E	1
252	A	6,1	F	2
253	A	6,2	G	3
254	A	6,3	H	4
255	A	6,4	I	5
00D	B	6,0	J	6
00E	B	6,1	K	7
00F	B	6,2	L	8
010	B	6,3	M	9
011	B	6,4	N	10

Now we're ready to perform the following:

1. Shutdown databases and/or applications using the disks to be migrated.
2. Unassign the volumes from the SCSI host.
   a. Go to **ESS Specialist**.
   b. Go to **Storage Allocation**.
   c. Go to **Open SYstem Storage**.
   d. Go to **Modify Volume Assignments**.
   e. Select the volumes from the list.
   f. Under Action, select **Unassign selected volume(s)**.

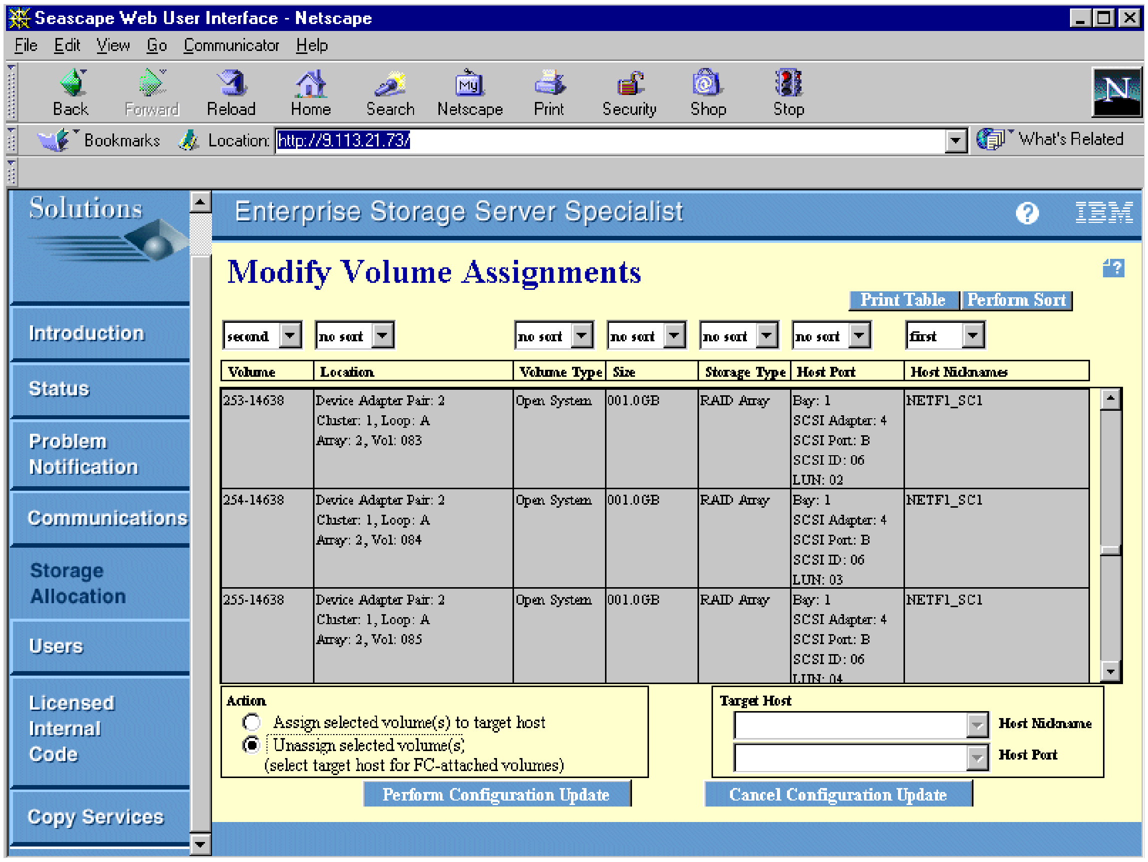

Figure 40. Modify Volume Assignments frame on ESS Specialist

g. Click on **Perform Configuration Update**. Once done, click **OK** (Figure 40).

3. Assign the volumes to the Fibre Channel host. See Figure 41.
   a. Go to **Modify Volume Assignments**.
   b. Select the volumes from the list.
   c. Under Action, select **Assign selected volume(s) to target host**.
   d. Select the Fibre Channel host adapter to which you want to assign the volumes.
   e. Click **Perform Configuration Update**. Once done, click **OK**.

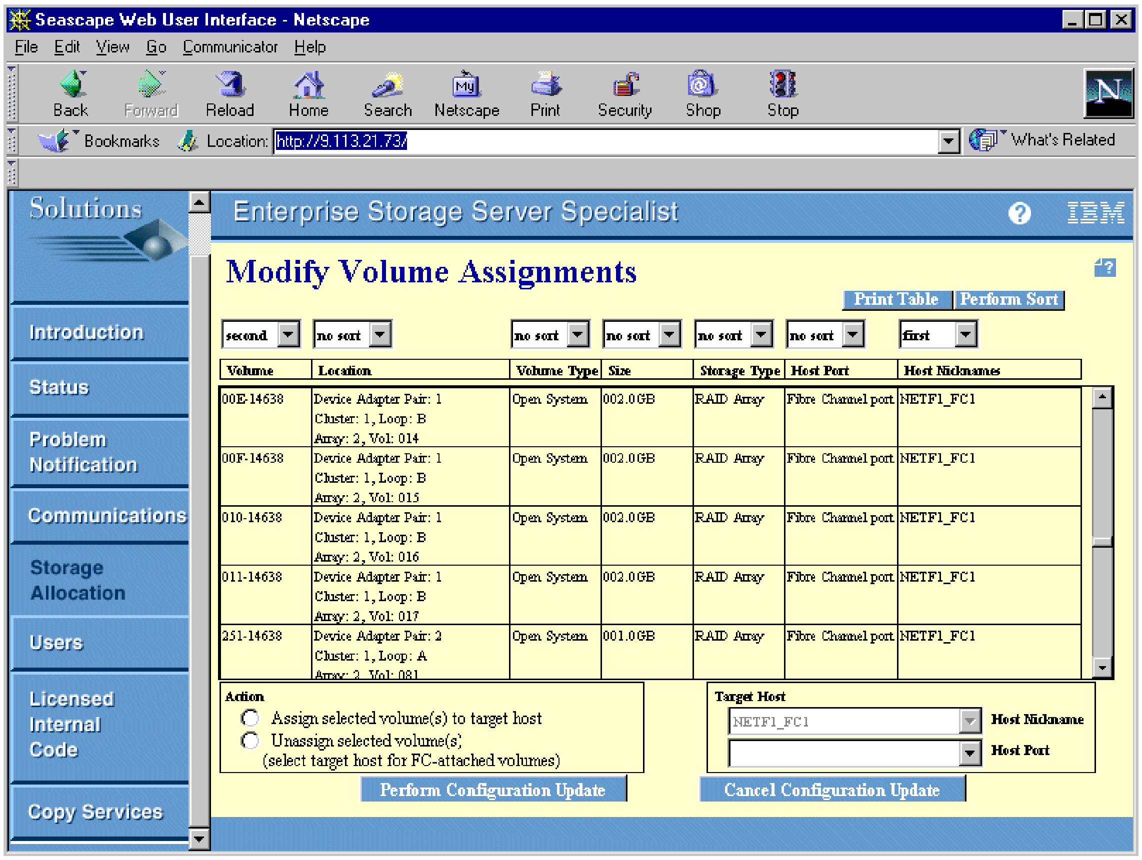

Figure 41. Volumes assigned to Fibre Channel host

   f. Reboot the host for the changes to take effect.

4. Open SAN Explorer. The view should look similar to Figure 42.

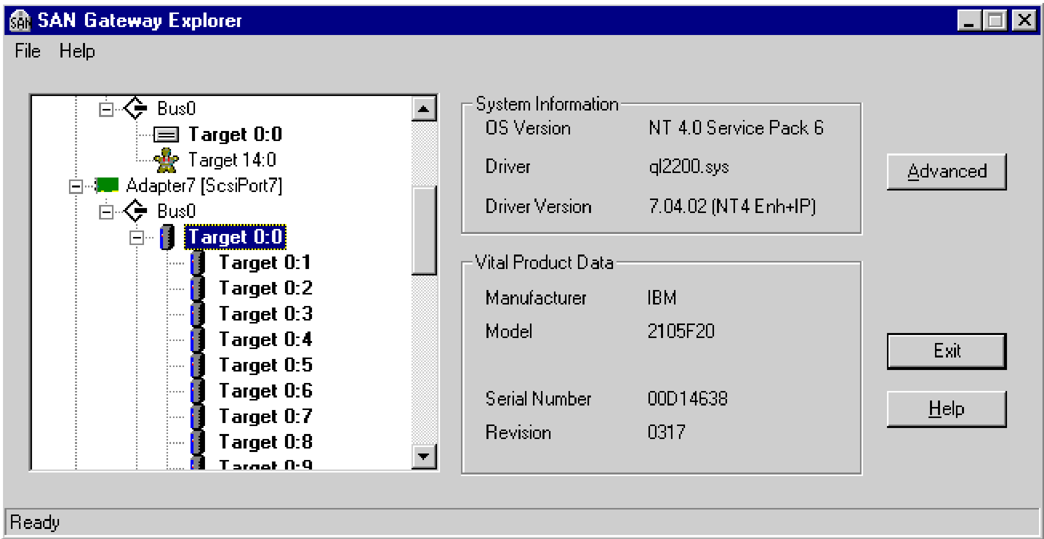

Figure 42. SAN Explorer view after migration

Table 5. Volume mapping after migration

Volume ID	SCSI Port on the ESS	SCSI Target ID & LUN	Fibre Channel Target & LUN Assignment	Disk Administrator Logical Drive Assignment	Disk Administrator Disk Number
251	A	6,0	0,5	E	6
252	A	6,1	0,6	F	7
253	A	6,2	0,7	G	8
254	A	6,3	0,8	H	9
255	A	6,4	0,9	I	10
00D	B	6,0	0,0	J	1
00E	B	6,1	0,1	K	2
00F	B	6,2	0,2	L	3
010	B	6,3	0,3	M	4
011	B	6,4	0,4	N	5

In Table 5, we assigned all the volumes (00D-011 and 251-255) to the Fibre Channel host at the same time. This changed the order of presentation to SAN Explorer and Disk Administrator, because LUNs 00D-011 will get lower LUNs than 251-255. ESS assigns LUNs based on the sequence of assigning, and the volume IDs. The earlier the LUNs are assigned and the lower their volume ID, the lower the LUNs they get.

This however should not affect the data on the disks. They will retain their logical drive assignments, their labels, and whatever files are on them before the migration.

The Disk Administrator view will be similar to the one shown in Figure 43.

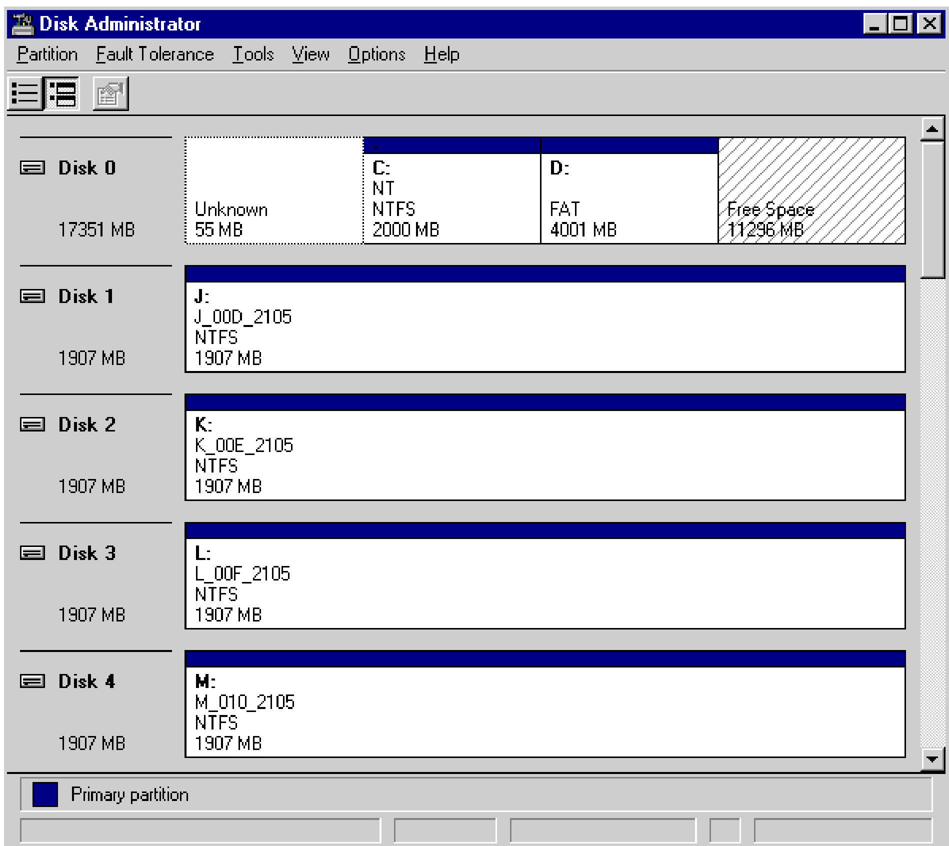

Figure 43. Disk Administrator view after migration

Take note that the internal disk got Disk number 0 (ahead of the ESS volumes), because the adapter to which it is connected is scanned first before the Fibre Channel HBA.

Here, drives J, K, L, M and N are shown ahead of E, F, G, H, and I.

If we want the order shown, where the order is E - N (alphabetically), then the steps are:

1. Assign first volumes 251-255.
2. Perform Configuration Update.
3. Assign next 00D-011.
4. Perform Configuration Update.

6.1.4 Migrating SCSI to FC on HP-UX server

This section describes the steps required on HP-UX to migrate disks/volumes on ESS configured from SCSI attached disks to the Fibre Channel attached disks.

1. Shutdown databases and/or applications using the disks to be migrated.
2. Unmount file systems (umount <file system>) with SAM (System Admin Menu utility).
3. Identify all disks assigned to the volume group affected on ESS unit with the ESS specialist.
4. Deactivate volume group via SAM.
5. Export the volume group. You need to enter a file name (full path) where you want to store logical volume information for the volume group being exported. (Note: We did not put a file name the first time since it was optional. When we tried to import the volume group we got the error message that the map file missing. Then we tried again and the final screen display looks as follows (Figure 44), before we selected **OK**.)

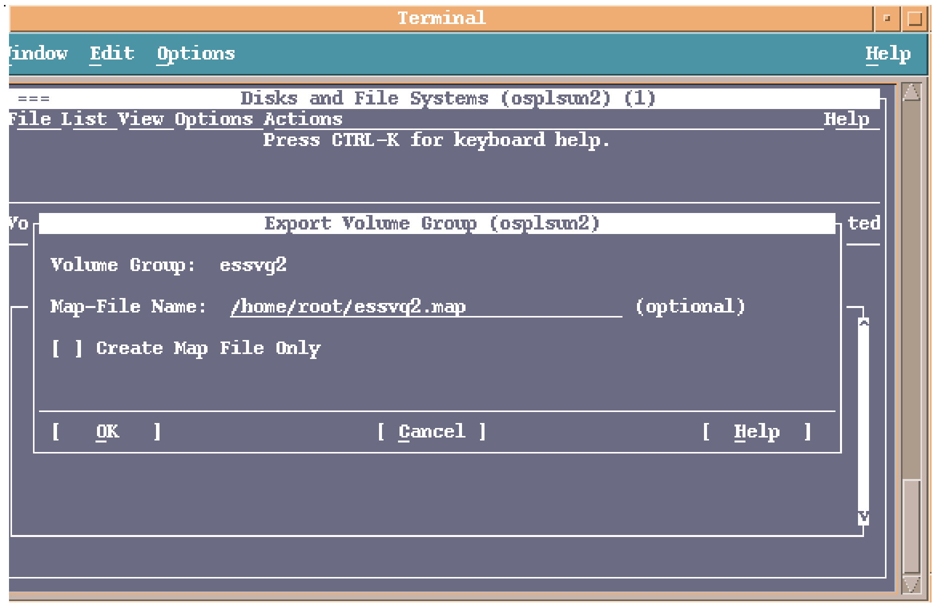

Figure 44. SAM's exportvg screen output

6. Now you can disconnect the SCSI cable and connect the fiber cable to the server and ESS unit. You do not have to remove the SCSI cable.
7. Unassign all LUNs on the ESS using the ESS Specialist which were assigned to the server by SCSI and assign them to a Fibre Channel port. This is where WWPN information becomes important, because this will be used to configure the ESS in the **Modify Volume Assignments** screen. See Figure 40 on page 59.
8. After successful configuration completion on ESS you should be able to see those disks next time you check disk devices.
9. Enter `sam` at your server command prompt. Select **Disks and File Systems** then press Enter.
10. Select **Volume Groups** then use the tab key to select the pull down menu item from the top.
11. Select **Actions** and press Enter. From listed option select **import** then press enter.
12. You will see all available volume groups can be imported. Select the one you exported and you will see all disks associated to that volume group.

13.Now use the tab key to go the next field called **new volume group name** and type in volume group name. You can use the same name or any name you like.

14.Then tab to the next field and type in the map file name which was created while exporting the volume group. See Figure 44.

15.Then use the tab key to select **OK** to import the volume group (Figure 45).

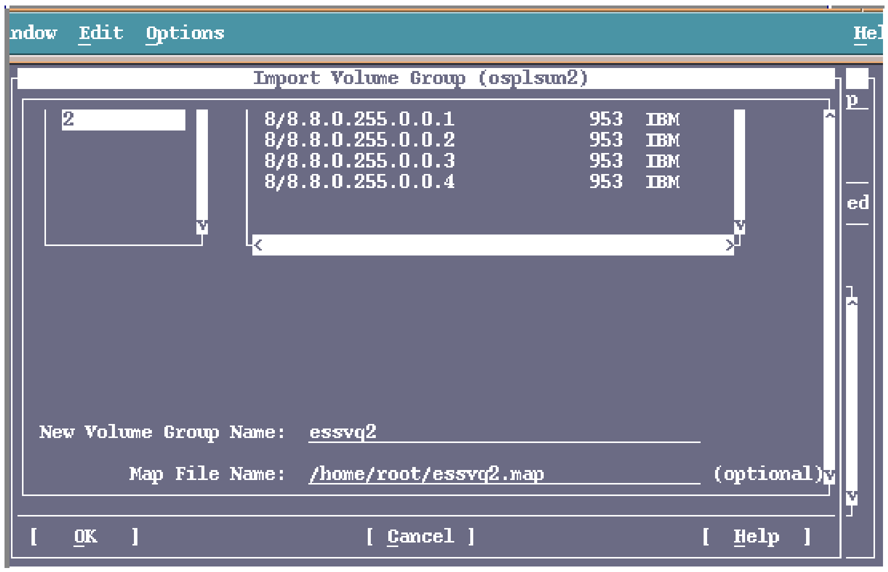

Figure 45. Screen display of SAM of import volume group

16.After successful command completion, you will now have active **essvg** volume group. You can mount all your file systems, which were unmounted in the beginning of this migration process.

17.Once all file systems are available you can start your databases and/or applications.

6.2 Concurrent migration

This section describes the process of migrating disks/volumes using a concurrent method. This method is concurrent from the operating system point of view. That is, the host does not lose any access to data while the migration is carried out.

The basic assumptions used here are:

- Your server has disks / logical volumes, that are located on an ESS in use. They are attached to the server with two or more SCSI interfaces.
- Subsystem Device Driver (SDD) is installed and running properly.
- The server has one or more FC adapters, the so called Host Bus Adapter (HBA) is installed, the ESS has FC adapter (Feature Code 3022) installed, and the interconnections between these has been established.
- All appropriate software prerequisites, drivers, and PTFs have been implemented.

The steps to achieve concurrent migration means:

1. Test that you have connectivity along both SCSI paths, see Figure 46.

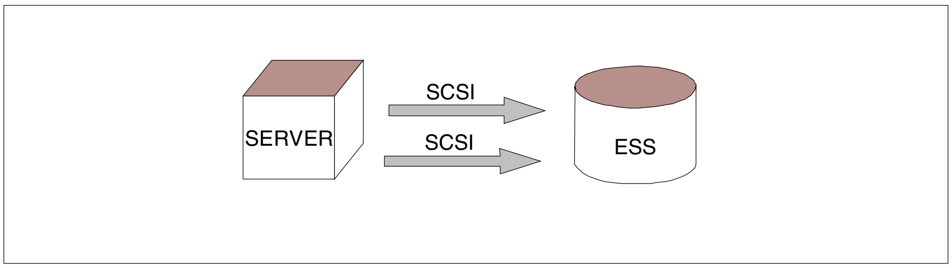

Figure 46. Multipath storage to server connection, SCSI only

2. Remove the disks from one of the SCSI ports on the ESS.
3. Attach these disks to the WWPN of one of the Fibre Channel HBAs; see Figure 47.

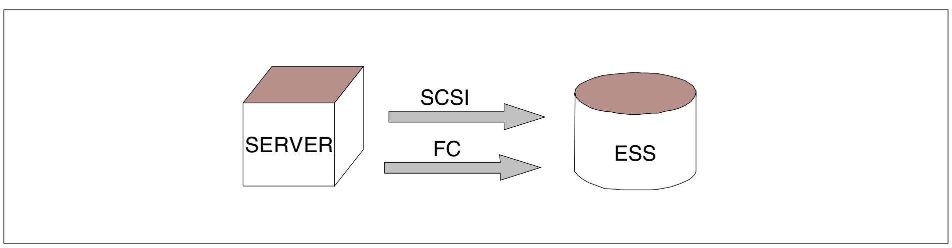

Figure 47. Multipath storage to server connection, SCSI and FC

4. Test the system to check that you have connectivity down both the SCSI and Fibre Channel paths.
5. Remove the disk attachment to the second SCSI port on the ESS and assign the disks to the WWPN of the second Fibre Channel HBA; see Figure 48.

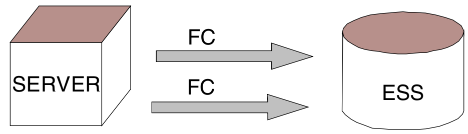

Figure 48. Multipath storage to server connection, FC only

6. Test the system to ensure that you have connectivity along both Fibre Channel paths.

6.2.1 Migration scenario with a Windows NT host

Here is the migration scenario staging the goal, environment, conditions, and procedures:

- **Goal:** To migrate concurrently two SCSI paths to Fibre Channel (Figure 49).
- **System environment:**

  Two SCSI adapters on the host side, each with a path going to a port on the same adapter at the ESS side.

  Two fibre adapters both on the host and ESS, configured but are not yet used.

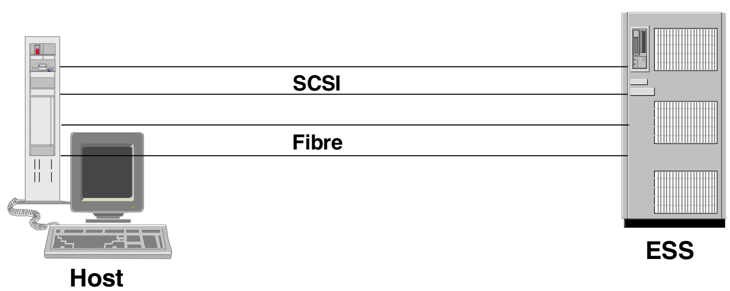

Figure 49. System setup for concurrent migration

The following conditions have been met before the migration:

- **Host side:**
  - FC HBAs are already installed on the appropriate host PCI slots.
  - Device drivers for the FC adapters are already installed and started.
  - SDD with FC support is already installed and started.
  - Two hosts (one for each FC adapter) are already defined on the ESS by using the ESS Specialist.
- **ESS side:**
  - Feature Code 3022 (Fibre Channel) adapters are already installed.
  - ESS Specialist is enabled and running successfully.
- **Others:**
  - Fiber links are already setup.
  - Switch is used to interconnect fibre host and ESS (this is optional).
  - Host and ESS are joined using port zoning on the switch.
  - Point to point protocols are used on both host and ESS adapters (although this is not necessary, but preferred if possible).
- **Procedure:**
  1. Go to **Start, Programs, Subsystem Device Driver, Subsystem Device Driver Management**.
  2. On the command prompt, type `datapath query adapter`. This shows the links available and their status (Figure 50).

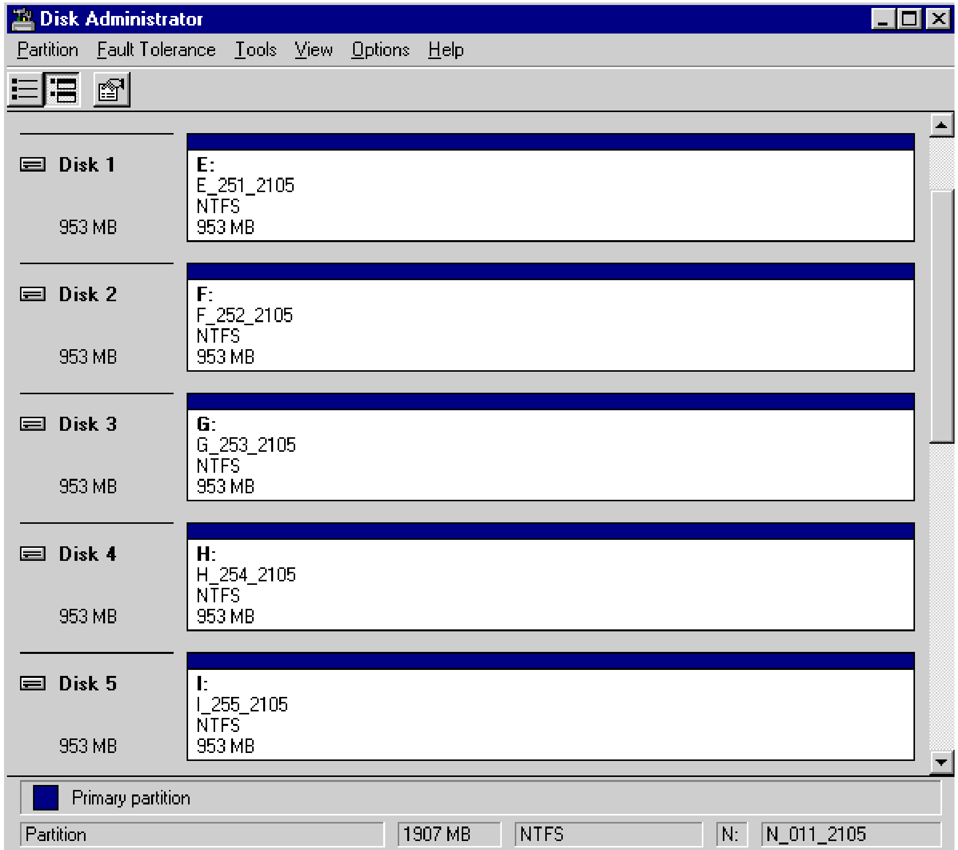

Figure 50. SDD command output showing two active SCSI paths

3. Run Disk Administrator to check that the host only sees the correct number of online and offline disks (in this case ten online, ten offline). See Figure 51.

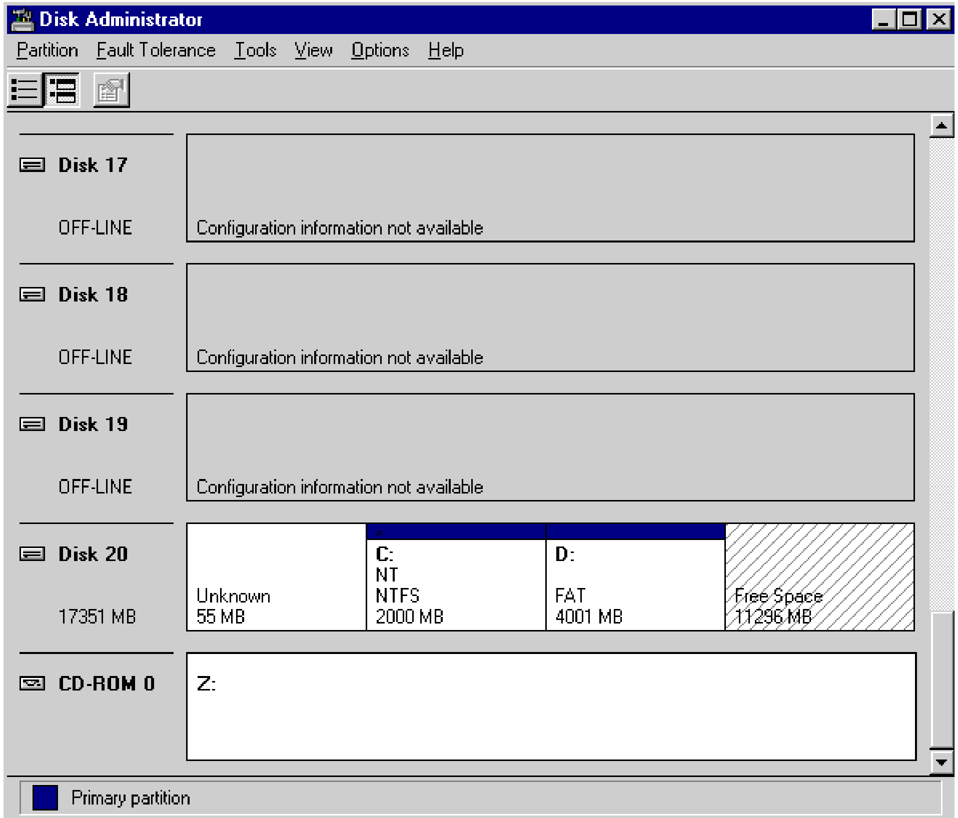

Figure 51. Shows how SDD handles presentation of two paths to 10 volumes

4. Run a disk I/O on a volume that is attached to multiSCSI paths. In this case, we use IOMETER. See Figure 52.

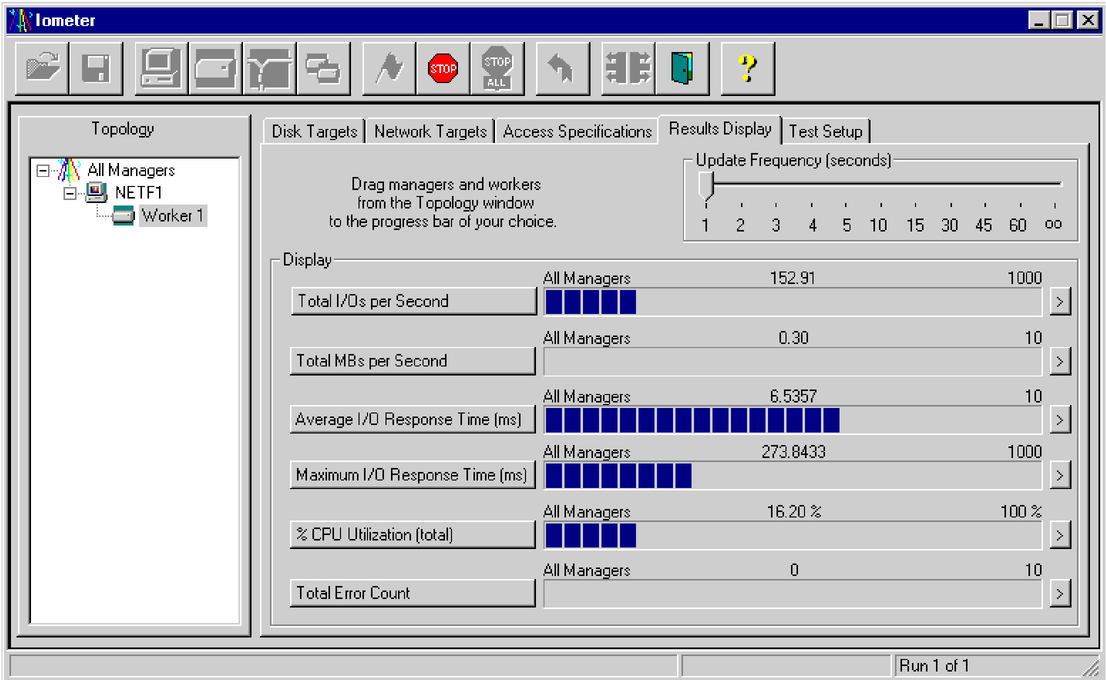

Figure 52. I/O test using Iometer

5. Go to ESS Specialist and assign the same volumes (which are already assigned to the two SCSI paths) to one of the fibre hosts. Refer to Figure 53.

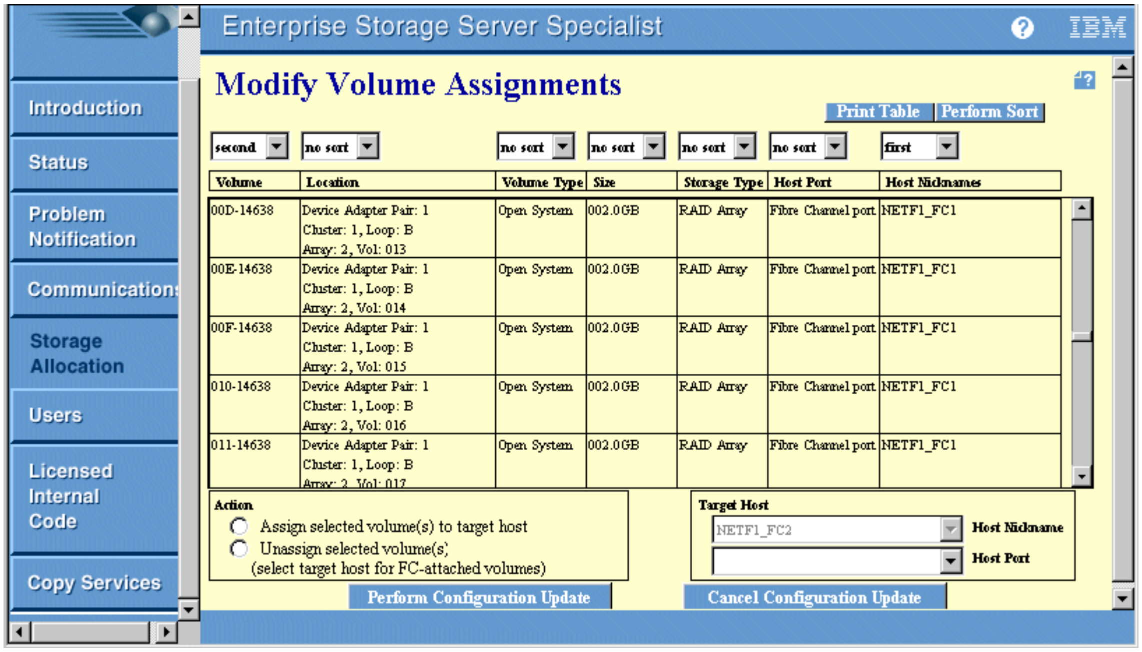

Figure 53. Volumes assigned to an additional host

6. Run Disk Administrator on the host. This enables the adapters (both SCSI and fibre) to re-scan the bus.You should get the same number of online but additional offline disks (in our example in Figure 54, twenty offline disks).

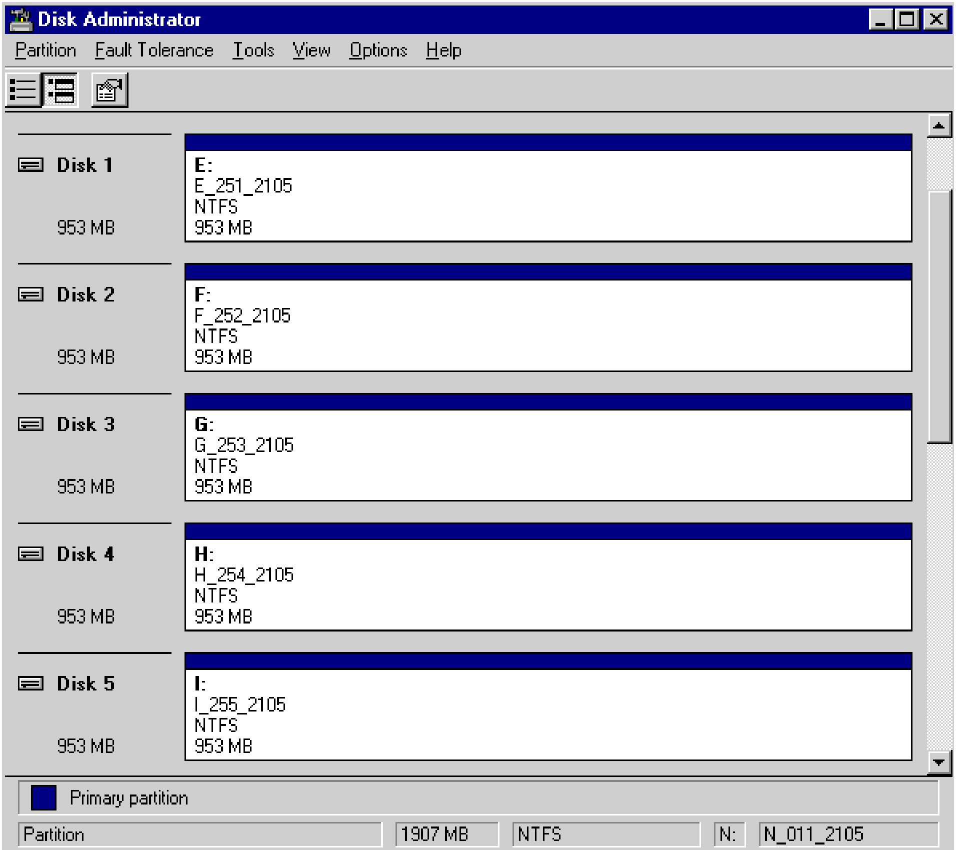

Figure 54. Shows how SDD handles presentation of three paths to 10 volumes

7. On the Subsystem Device Driver Management window, type `datapath query adapter`. This now shows the first fibre link that we just added. Notice Adapter 2, SCSI Port 7. This is the new path that we just added via host NETF1_FC1 (Figure 55).

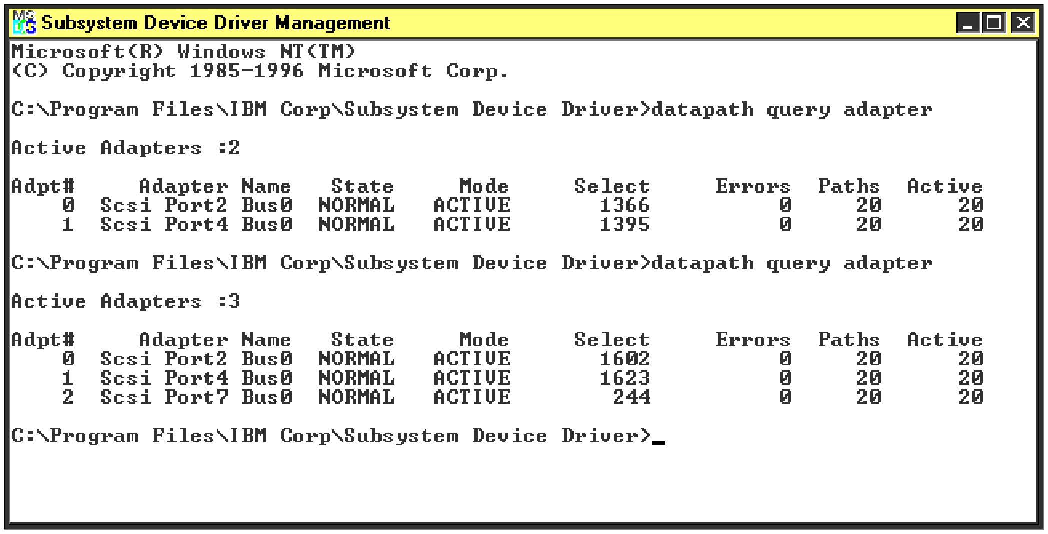

```
Subsystem Device Driver Management
Microsoft(R) Windows NT(TM)
(C) Copyright 1985-1996 Microsoft Corp.

C:\Program Files\IBM Corp\Subsystem Device Driver>datapath query adapter

Active Adapters :2

Adpt#     Adapter Name   State     Mode       Select      Errors  Paths  Active
    0   Scsi Port2 Bus0  NORMAL    ACTIVE       1366           0     20      20
    1   Scsi Port4 Bus0  NORMAL    ACTIVE       1395           0     20      20

C:\Program Files\IBM Corp\Subsystem Device Driver>datapath query adapter

Active Adapters :3

Adpt#     Adapter Name   State     Mode       Select      Errors  Paths  Active
    0   Scsi Port2 Bus0  NORMAL    ACTIVE       1602           0     20      20
    1   Scsi Port4 Bus0  NORMAL    ACTIVE       1623           0     20      20
    2   Scsi Port7 Bus0  NORMAL    ACTIVE        244           0     20      20

C:\Program Files\IBM Corp\Subsystem Device Driver>_
```

Figure 55. Datapath query adapter output with additional link

8. Go to ESS Specialist and assign the same volumes (which are already assigned to the two SCSI and one fibre paths) to the other fibre host (Figure 56).

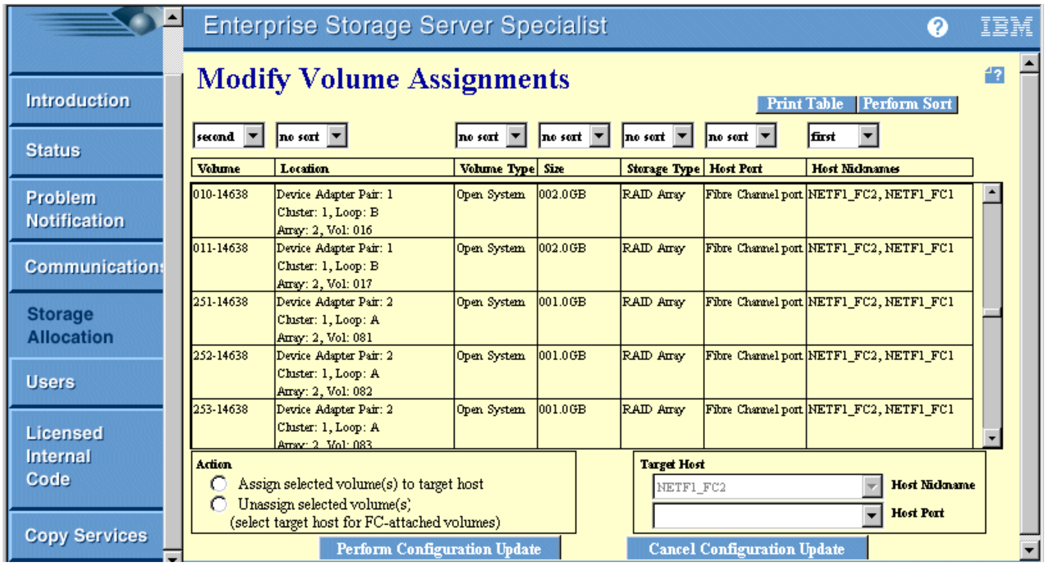

Figure 56. Volumes assigned to two fiber paths

9. Run Disk Administrator again on the host to re-scan the bus. You should get the same number of online but more additional offline disks (in our example in Figure 57, thirty offline disks).

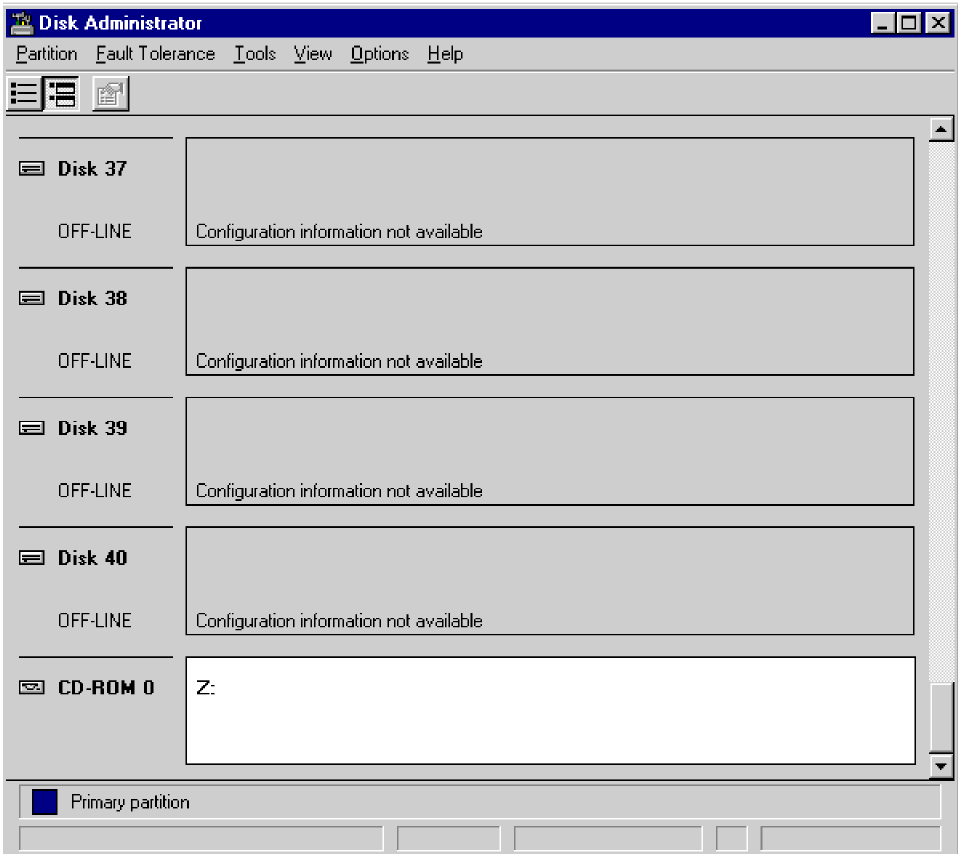

Figure 57. Shows how SDD handles presentation of four paths to 10 volumes

10. On the Subsystem Device Driver Management window, type `datapath query adapter`. This now shows the second FC link that we just added. Notice Adapter 3, SCSI Port 8. This is the new path that we just added via host NETF1_FC2. See Figure 58.

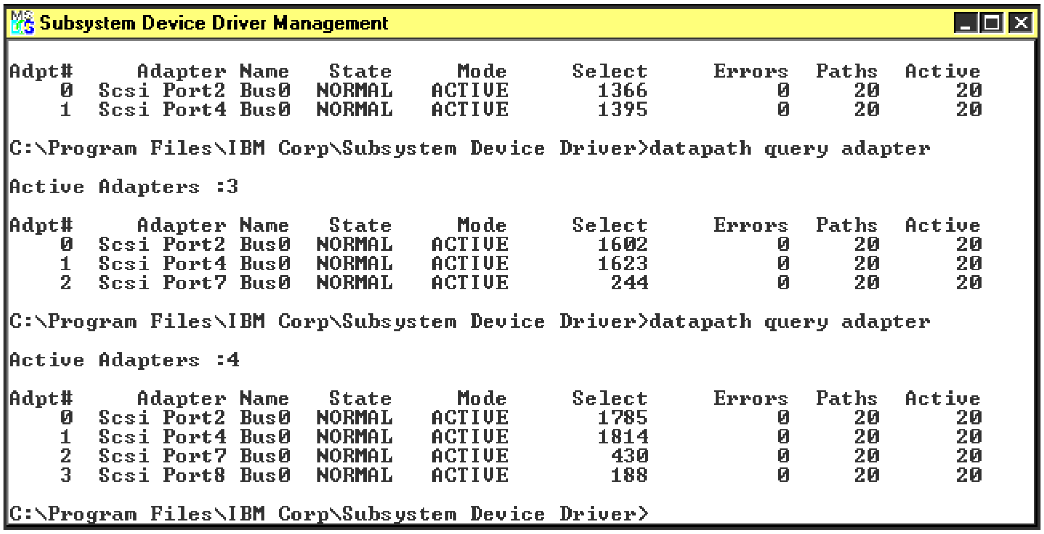

```
Subsystem Device Driver Management

Adpt#     Adapter Name   State    Mode      Select     Errors  Paths  Active
    0   Scsi Port2 Bus0  NORMAL   ACTIVE      1366          0     20      20
    1   Scsi Port4 Bus0  NORMAL   ACTIVE      1395          0     20      20

C:\Program Files\IBM Corp\Subsystem Device Driver>datapath query adapter

Active Adapters :3

Adpt#     Adapter Name   State    Mode      Select     Errors  Paths  Active
    0   Scsi Port2 Bus0  NORMAL   ACTIVE      1602          0     20      20
    1   Scsi Port4 Bus0  NORMAL   ACTIVE      1623          0     20      20
    2   Scsi Port7 Bus0  NORMAL   ACTIVE       244          0     20      20

C:\Program Files\IBM Corp\Subsystem Device Driver>datapath query adapter

Active Adapters :4

Adpt#     Adapter Name   State    Mode      Select     Errors  Paths  Active
    0   Scsi Port2 Bus0  NORMAL   ACTIVE      1785          0     20      20
    1   Scsi Port4 Bus0  NORMAL   ACTIVE      1814          0     20      20
    2   Scsi Port7 Bus0  NORMAL   ACTIVE       430          0     20      20
    3   Scsi Port8 Bus0  NORMAL   ACTIVE       188          0     20      20

C:\Program Files\IBM Corp\Subsystem Device Driver>
```

Figure 58. Datapath query adapter output with another additional link

11. Go to ESS Specialist and unassign the volumes from the SCSI ports. Refer to unassigning volumes, list item 3 on page 52.

12. After performing the configuration update, go to Subsystem Device Driver Management window, and type `datapath query adapter`. It may take some time before SDD tags the link as Degraded. Keep on refreshing the window by executing the same `datapath query adapter` command. After a while, the State of the SCSI links should become *Degraded as* shown in Figure 59.

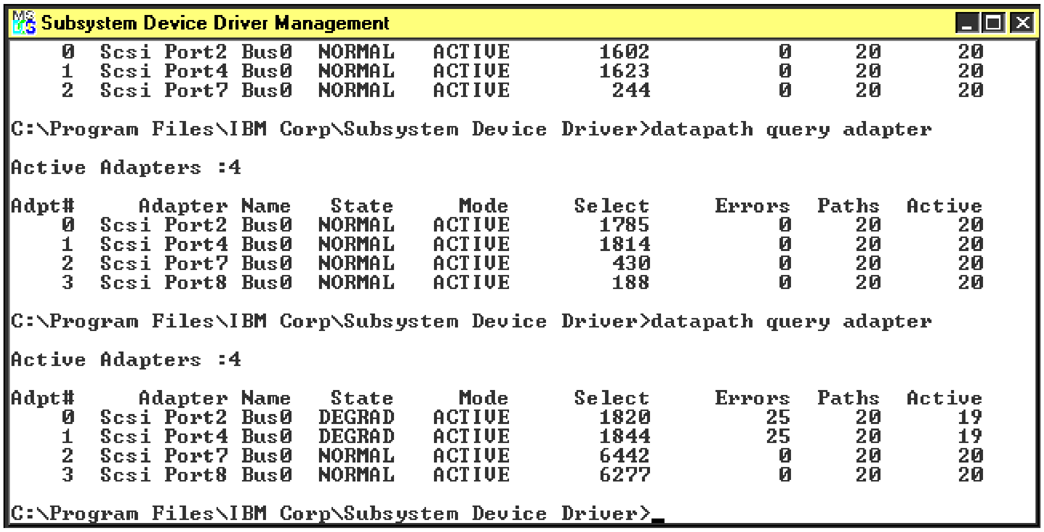

Figure 59. SDD command output after unassigning volumes to the SCSI paths

You should not get an error from your I/O application, because the two FC paths are already active. See Figure 60.

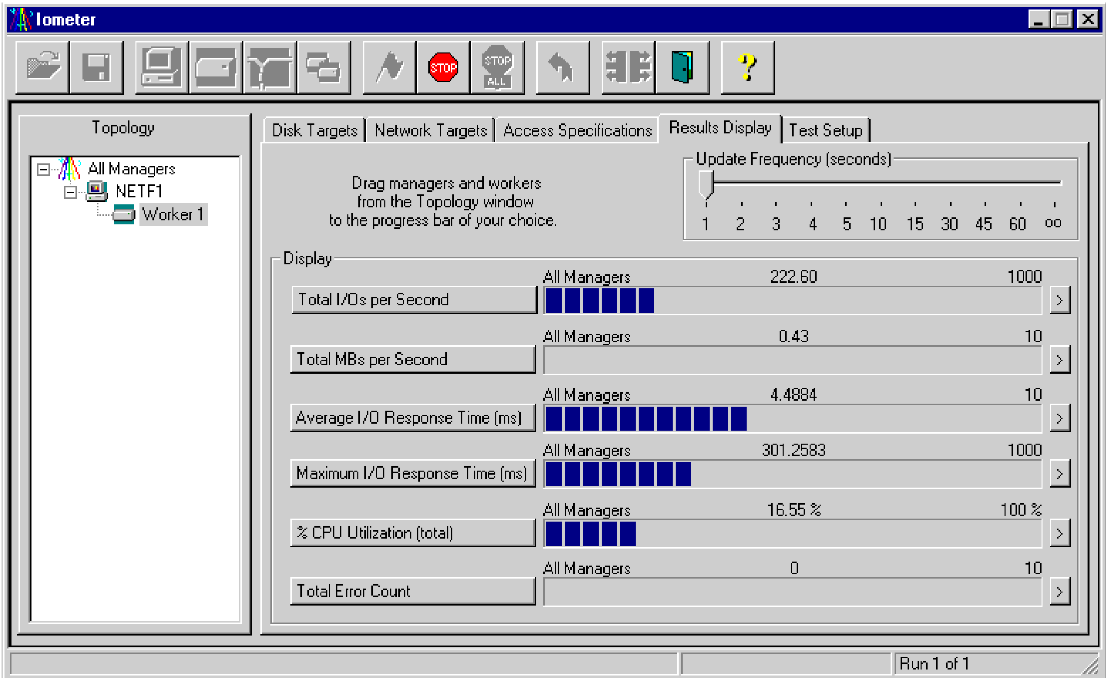

Figure 60. Iometer still runs after the SCSI paths have been removed

The two inactive paths will still be seen and checked by SDD, because SDD is not aware that it was removed manually. However, this shouldn't affect your capability to access the data through the two FC paths. Until you reboot the host, the two SCSI paths will still appear on the SDD datapath adapter query.

Chapter 7. Migrating from SAN Data Gateway to Fibre Channel

In SAN implementations there is a need to use bridges or SAN Data Gateways (SDG) to connect between the fibre channel and SCSI, SSA, and other interfaces on storage devices which do not have native fibre channel support. This allows you to utilize your existing storage, while having the benefits of Fibre Channel.

However, the need for migration to end-to-end Fiber Channel becomes inevitable, driven by the need for faster, larger, and more reliable storage requirements. This chapter provides details on how to migrate a data gateway based SAN to pure Fibre Channel, including the components, considerations, and impacts.

7.1 The IBM SAN Data Gateway (SDG)

The SAN Data Gateway is basically a protocol converter between FCP and SCSI, and supports tape drives (MP3570 & Magstar 3590), tape libraries (Magstar MP3575, Magstar 3494 & IBM 3502 DLT) and disk (ESS) subsystems on Intel based servers running Windows NT and Unix based machines. For the latest list of supported operating systems & storage products, visit the IBM SAN Data Gateway Web site at:

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

The SDG is a fully scalable product with up to three Fibre Channel and four Ultra SCSI Differential interfaces for disk and tape storage attachment. Each Fibre Channel interface could support dual or single shortwave ports (for a maximum of six shortwave ports), and single longwave ports (for a maximum of two longwave ports). Below are the details of the interfaces on a SDG:

- Fibre Channel
  - Supports both loop (private and public) and point to point topologies.
  - Connection to node or switch could go up to 500m (50 micron) for shortwave and 10km for longwave.
- SCSI
  - SCSI channels have automatic speed negotiation capability for wide or narrow bus widths and Standard, Fast or Ultra speeds.
  - Each channel supports up to 15 target IDs and up to 32 LUNs per ID (subject to an overall total of 255 devices).
  - Cable lengths can go up to 25m.

Note: Although the SDG has a limitation of 255 LUNs (actually 256, but one LUN is used by the SDG for command and control), you still need to consider the limitations on your host operating system, adapter and storage subsystem. Whichever is the lowest dictates the maximum that you could use. For example, you have the following configuration shown in Table 6.

Table 6. Operating system LUN limitation

System Component	LUN Limit
Windows NT 4.0 w/ SP3	120 (15 target IDs x 8 LUNs)
Host Bus Adapter (HBA)	256
ESS	960 (15 target IDs x 64 LUNs)
SDG	255

In the above example, you will have a maximum LUN limit of 120, since it is the lowest supported by one of the components of your system, which is the operating system. However, fixes could be available to eliminate some problems and limitations. SP4 and above already supports up to 256 devices.

- Ethernet
  - 10BaseT port for out-band management (using StorWatch SAN Data Gateway Specialist)
- Service Port
  - 9 pin D-shell connector for local service, configuration and diagnostics

Also, the SDG allows *channel zoning*, as well as *LUN masking* between its FC and SCSI ports. These enable you to specify which Fibre Channel hosts could connect to LUNs defined on which SCSI ports. This is essential if you don’t want multiple hosts accessing the same LUNs, to avoid data integrity problems.

For more information on the IBM SDG, refer to *IBM Storage Area Network Data Gateway Installation and User’s Guide*, SC26-7304, or the redpiece *Planning and Implementing an IBM SAN*, SG24-6116.

7.2 Migration considerations and steps

Similar to the native SCSI to FC migration, SDG to FC migration requires these steps. The general steps (on the ESS side) involved in migrating volumes from SDG to native FC are:

1. Define the new FC HBA to the ESS using the ESS Specialist.
2. Unassign the volumes from the SCSI host.
3. Assign the volumes to the newly defined host. In practice this means assigning the volumes to the WWPN of the host adapter.

This assumes that there is a spare slot in the ESS host bays for the FC adapter to coexist with the SCSI adapter. If the SCSI adapter has to be removed before the FC adapter can be installed, care must be taken to list the volume assignments before the SCSI adapter is removed.

Important

Removing a SCSI adapter from an ESS host bay causes all the volume assignments for that adapter to be lost.

7.2.1 Migration example for Windows NT 4.0

Here is an example of a detailed migration from SDG to FC on Windows NT 4.0. The system environment is as follows:

- **Number of ESS volumes defined for the NT host:** 10
- **SDG connection:** 5 volumes each on SCSI channel 1 & 2; FC host connected to SAN connection 1; channel zoning used to connect SAN connection 1 to SCSI channel 1 & 2.

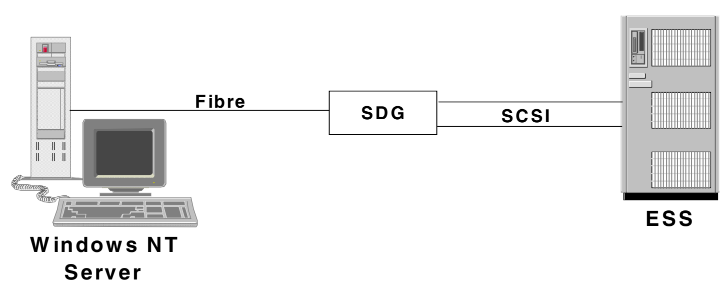

Figure 61. System setup for SDG to FC migration

- **Software utilities used**:
  - Browser (Netscape or Internet Explorer)
  - SAN Data Gateway Specialist
  - SAN Explorer
  - Disk Administrator

SAN Data Gateway Specialist and SAN Explorer can be downloaded at:

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

Here's how the initial setup looks on the SAN Data Gateway Specialist (Figure 62).

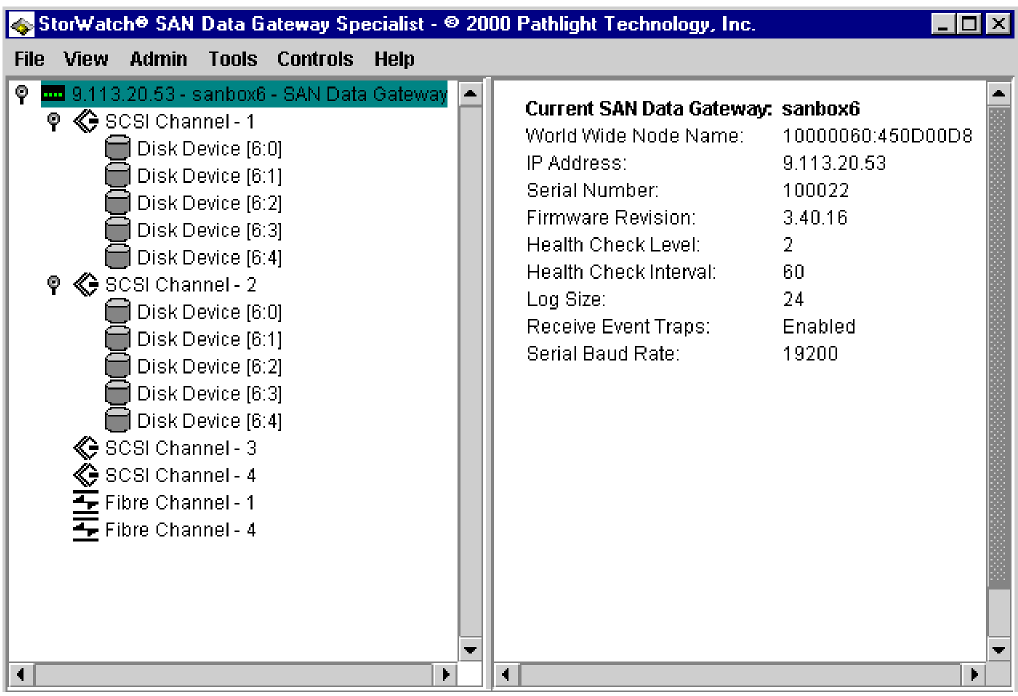

Figure 62. SDG Specialist view of the volumes

Using SAN Explorer, this is how the disks are presented by the SDG to the FC host (see Figure 63).

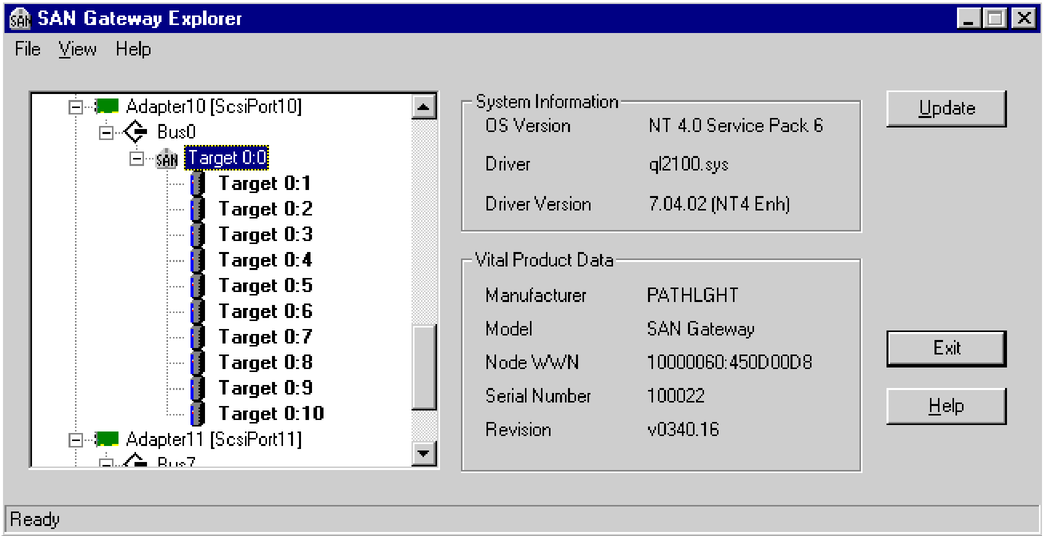

Figure 63. Fibre host view of the disks using SAN Gateway Explorer

The volumes already have labels and logical drive assignments. This is how the volumes look in Disk Administrator (see Figure 64).

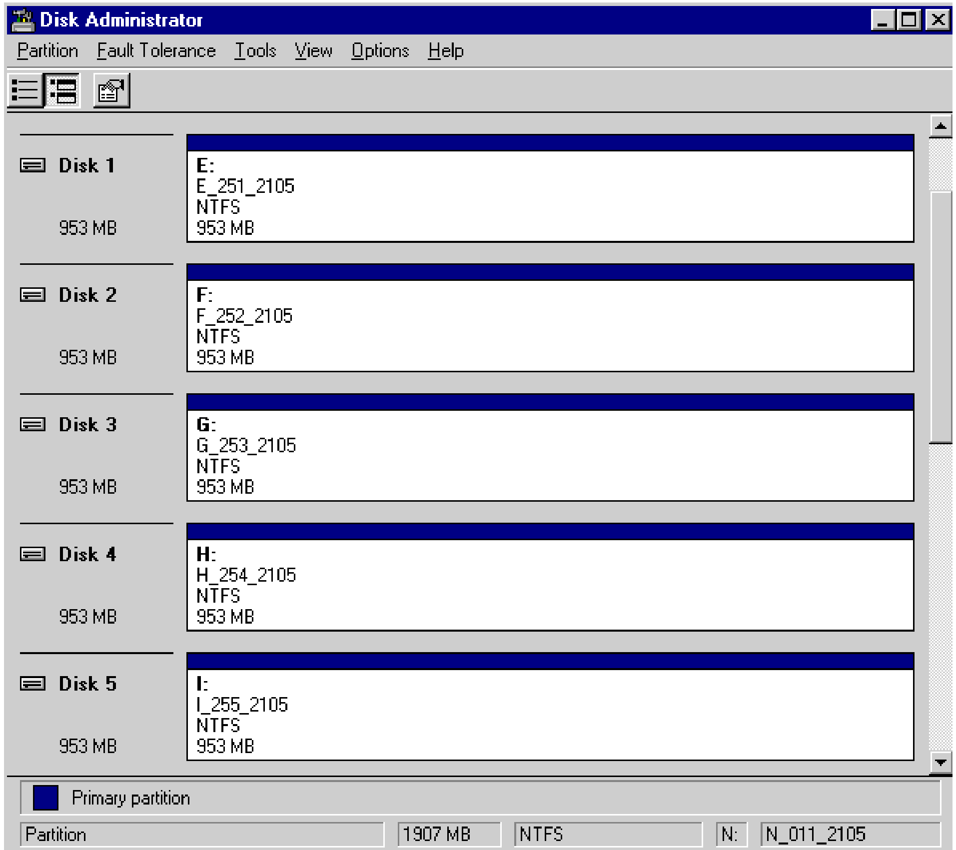

Figure 64. Disk Administrator view of the ESS disks

Take note that only five volumes (out of ten) are shown because of the limitation of the Disk Administrator window.

Table 7 shows how the volumes are mapped on the ESS, the SDG, and the host.

Table 7. Relationship between SCSI ID and disk number

Volume ID	SCSI Channel on the SDG	SCSI Target ID & LUN	Fibre Target & LUN assigned by the SDG	Disk Administrator Logical Drive Assignment	Disk Administrator Disk Number
251	1	6,0	0,1	E	1
252	1	6,1	0,2	F	2
253	1	6,2	0,3	G	3
254	1	6,3	0,4	H	4
255	1	6,4	0,5	I	5
00D	2	6,0	0,6	J	6
00E	2	6,1	0,7	K	7
00F	2	6,2	0,8	L	8
010	2	6,3	0,9	M	9
011	2	6,4	0,10	N	10

Notice that the FC LUN assignment of the volumes is provided by the SDG based on SCSI channel (SDG side) and SCSI target ID & LUN. The lower the SCSI channel and target ID & LUN, the lower the FC LUN assigned.

Now, we're ready to perform the migration.

7.2.1.1 Define the new Fibre Channel host using ESS Specialist

Use these steps to define the new Fibre Channel host using ESS Specialist:

1. Log on to the ESS Specialist.
2. Go to **Open System Storage**.
3. Select **Modify Host Systems**.
4. Under Nickname, type in the nickname you want to identify the host (in our example, we use NETF1_FC1).
5. Under Host Type, select **PC Server (Win NT4.0 or higher)**.
6. Under Host Attachment, choose **Fibre Channel attached**.
7. Under World-Wide Port-Name, type in or select using the pull-down arrow the WWPN of the HBA on the host.
8. Under Hostname/IP Address, type in the hostname you want to identify the host (in our example, we use NF1_FC1).

9. Click ***Add>>***. The host should now show on the Host Systems List. An example is shown in Figure 65.

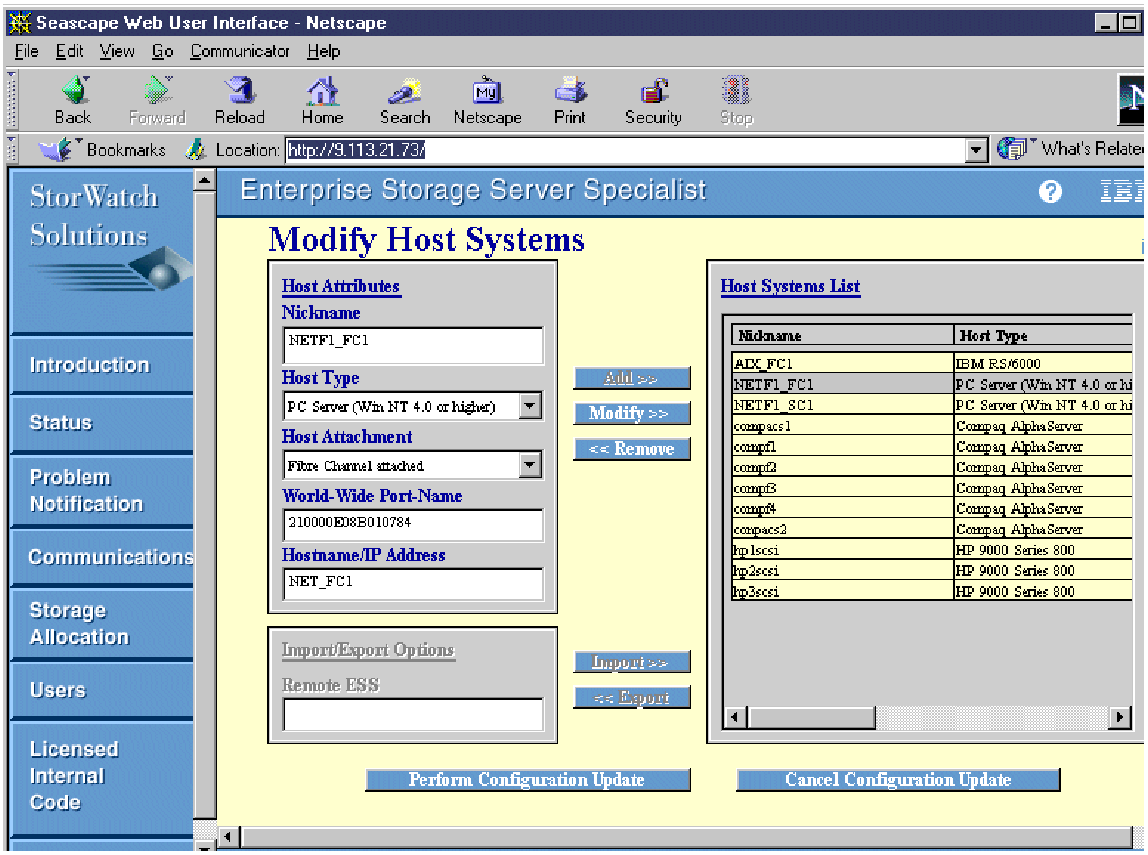

Figure 65. Defining a new host on ESS Specialist

10.Click **Perform Configuration Update**. Once done, click **OK**.

7.2.1.2 Unassign the volumes from the SCSI host

Use these steps to unassign the volumes from the SCSI host (Figure 66):

1. Go to **Modify Volume Assignments**.
2. Select the volumes from the list.
3. Under Action, select **Unassign selected volume(s)**.

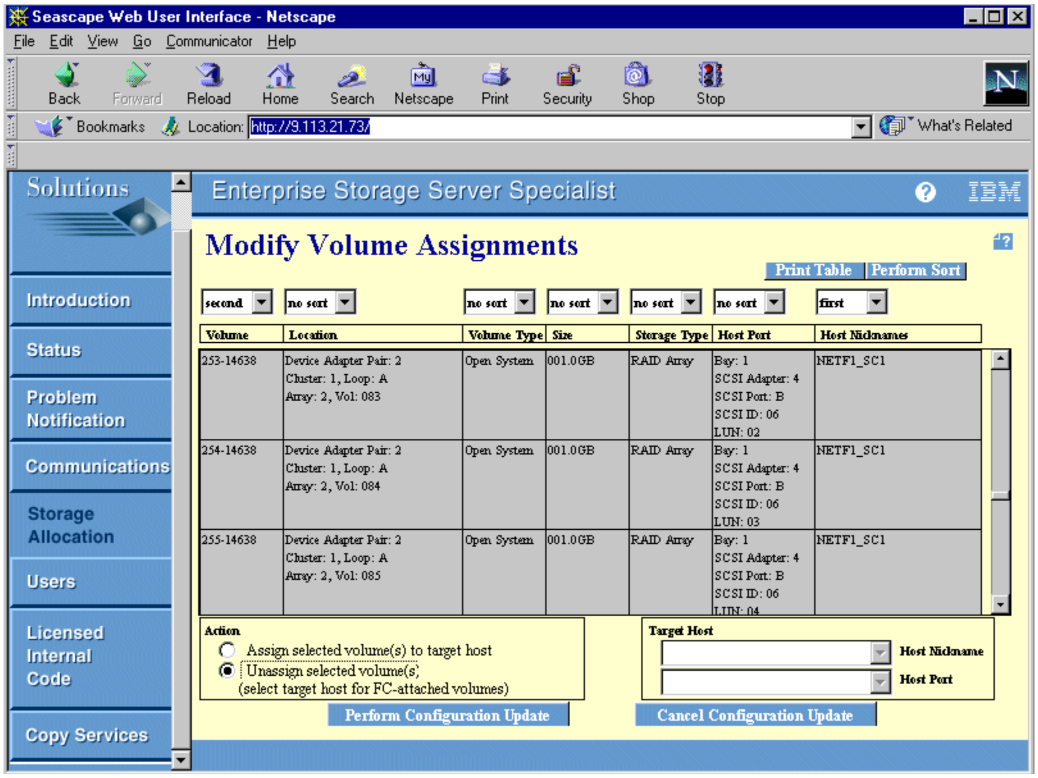

Figure 66. Unassigning the volumes from the SCSI host

4. Click **Perform Configuration Update**. Once done, click **OK**.

7.2.1.3 Assign the volumes to the Fibre Channel host

Use these steps to assign the volumes to the Fibre Channel host (Figure 67):

1. Go to **Modify Volume Assignments**.
2. Select the volumes from the list.
3. Under Action, select **Assign selected volume(s) to target host**.
4. Select the Fibre Channel host to which you want to assign the volumes.
5. Click **Perform Configuration Update**. Once done, click **OK**.

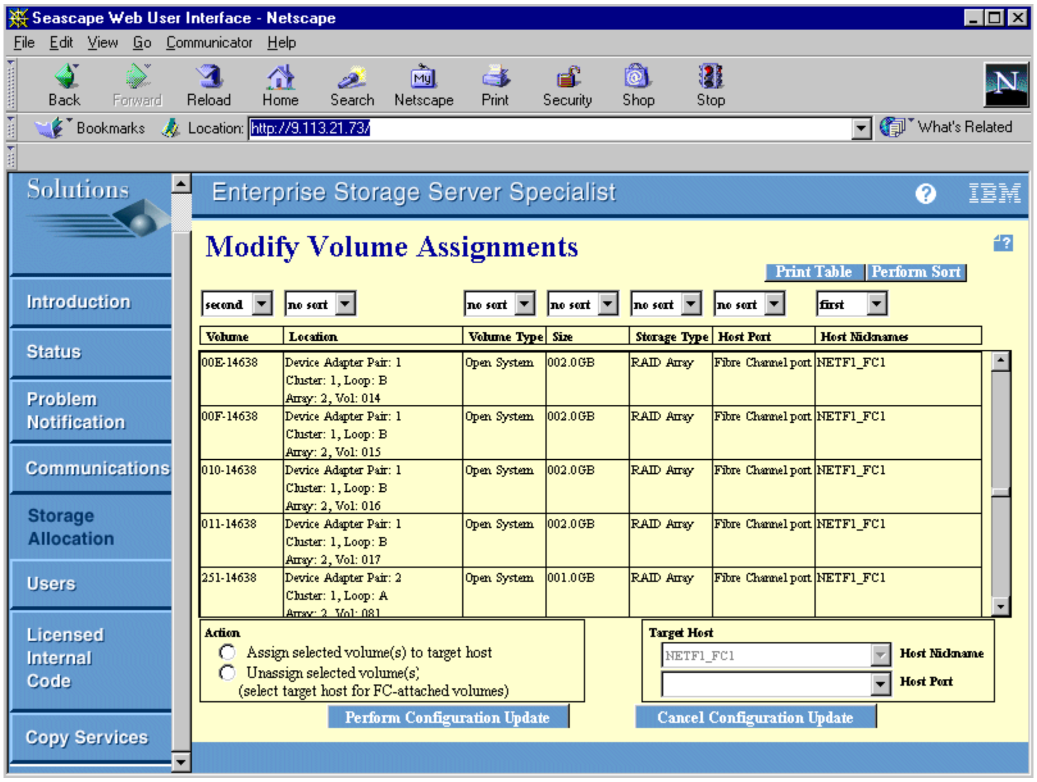

Figure 67. Volumes assigned to Fibre Channel host

6. Reboot the host for the changes to take effect.

Now that the volumes are direct FC attached, SAN Explorer and Disk Administrator should present the same disk volumes in the same way.

Table 8 compares the effect of the migration to the LUN assignments.

Table 8. Relationship between FC LUN ID and disk number

Volume ID	**Fibre Target & LUN assigned by the SDG (old)**	**Fibre Target & LUN assigned by the ESS(new)**	**Disk Administrator Logical Drive Assignment**	**New Disk Administrator Number**
251	0,1	0,5	E	6
252	0,2	0,6	F	7
253	0,3	0,7	G	8
254	0,4	0,8	H	9
255	0,5	0,9	I	10
00D	0,6	0,0	J	1
00E	0,7	0,1	K	2
00F	0,8	0,2	L	3
010	0,9	0,3	M	4
011	0,10	0,4	N	5

Notice that the new FC LUN numbering of the volumes moved one unit below (first LUN now starts at 0) because of the absence of the SDG.

The same as in SCSI migration, the LUN and the disk numbers may have changed, but they shouldn't affect what is written on the disks.

So if you want to retain the same disk numbering on Disk Administrator, you need to perform the following steps in sequence:

1. Assign first volumes 251-255.
2. Perform configuration update.
3. Assign volumes 00D-011.
4. Perform Configuration update.

Chapter 8. Performance considerations

This chapter provides information on aspects you should consider while thinking about what performance you can get from the new Fibre Channel attachment.

Fibre Channel as well as ESS offers a large range of possible configurations. We compare SCSI to Fibre Channel on the same configuration. Detailed information on ESS performance can be found in *IBM Enterprise Storage Server Performance Monitoring and Tuning Guide*, SG24-5656.

Table 9. ESS performance numbers

Workload	Throughput
Sequential read from disk - 64KB block size	386 MB/sec
Sequential writes to disk - 64KB block size	160 MB/sec
100% cache hit reads - 4KB blocks	45,000 SIO/sec
Random reads, no cache hits - 4KB blocks	12,000 SIO/sec
Random read/write, 70/30 ratio, 50% cache hits	15,000 SIO/sec
Random writes to disk - 4 KB blocks	3,800 SIO/sec
Writes 100% cache hits - 4KB blocks	12,000 SIO/sec

Table 9 shows the performance that has been obtained from the ESS in the performance testing laboratory in San Jose.

8.1 Bandwidth

The most obvious and highly touted performance characteristic of Fibre Channel is its bandwidth capability of 100 MB/sec full duplex, for a total of 200 MB/sec. This is far superior to the standard SCSI interfaces being used today, where 20, 40, 80 and 160 MB/sec implementations are available. Most commonly used are the 20 and 40 MB/sec implementations. However, there are a few other points that need consideration.

The bandwidth numbers cited above are maximum bandwidth capabilities of the architecture. A single physical link is able to transfer this bandwidth from the technical point of view. However, you should not plan workload based on these maximums.

In addition, the throughputs will vary based upon type of application, specific adapter and host attachment. As a general rule-of-thumb, one should base planning decisions assuming a throughput capability of around 60 MB/sec per Fibre Channel port configured to an ESS.

8.2 Fibre Channel attachment

While planning multipath Fibre Channel attachment via a SAN fabric you should always configure your paths for maximum throughput and avoid creating bottlenecks within the fabric by using too few paths. The guidelines are that a fibre channel path should be sized for a maximum throughput of 60MB/sec and 400GB of storage per port. It must also be understood that there may be occasions when not all the paths will be used.

A fabric can be built with only one switch. The Fibre Channel architecture also allows multiple switches to be interconnected, this is called switch cascading. When switches are cascaded, multiple paths between a host system and the storage device will be created. However, it is important to note that only the shortest paths between switches will be used. Both the Brocade and McData switches use a protocol called "Find Shortest Path First" (known as FSPF). This calculates the shortest path between a host and a device during fabric initialization and from then on, only that path and others equal in length to it will be used. If there are alternative routes through other switches which are longer they will not be used until all the shortest paths have failed.

The use of inter-switch links therefore needs careful consideration. Links between switches are desirable so that a common zoning configuration can exist across the fabric, it also means switches can be managed using IP over Fibre from only one switch ethernet connection.

Multiple links between switches only make sense if all the route lengths between host and storage are the same.

8.3 Number of ports required on ESS

Once the performance requirements of your application is well understood, you can use the guidelines just presented in the previous section to determine the proper number of host ports.

For example, if your application requires 100 MB/sec of throughput, then you would most likely need two Fibre Channel ports on the ESS to support this requirement. The reason for this is that although the rated speed of fibre channel is 100 MB/sec it is bad practice to size a component at its maximum capacity.

Sometimes you have little or no information about the exact performance requirements of the application. In these cases, you must make guesses on the number of ports needed. The following rules-of-thumb can help you in these cases:

1. Eight fibre connections within an ESS should handle the performance needs for nearly all applications. In this case, two fibre adapters should be installed in each of the four host bays.
2. Consider using four ports, one in each of the four host bays if you have:
   a. Primarily an OLTP workload in which you do not need the full sequential performance capability of the ESS.
   b. Small or moderate-capacity configurations, such as 1.6TB or less of open system fibre attached storage.
3. In some cases, you may not be planning to use the full subsystem capacity for accessibility from Fibre Channel hosts. For example, you may be mixing SCSI and ESCON attachments with your Fibre Channel attachments. In this case, it may be useful to plan the number of Fibre Channel ports based on the amount of disk capacity you plan to access through the Fibre Channel ports. A rule-of-thumb for planning is to use one Fibre Channel port for every 400 GB of data, with a two-port minimum to ensure data availability in the event of an interface failure.

8.4 Port positions in ESS

The Fibre Channel adapters are plugged into the four host bays in ESS. Two of these host bays are located in Cluster 1 and the other two are located in Cluster 2. These clusters are on separate power boundaries. Each host bay is composed of a single bus that interconnects the adapters on that bay. When populating the host bays with Fibre Channel adapters, you want to spread them as evenly as possible across all of the host bays in order to get the best performance. To get maximum performance on two FC interfaces attached to the same or different servers, you should install them in Bay 1 and Bay 4.

8.5 Loop versus fabric

In general, you obtain better performance with direct connection or fabric connection, and that sharing bandwidth with FC-AL limits performance for the hosts sharing the loop. Some hosts, RS6000, for example, when directly connected to the ESS give better performance when using the FC-AL protocol rather than Point-to-point protocol.

Appendix A. ESS Open Systems support

ESS is supported in a wide variety of environments which is being increased. Manufacturers are constantly changing operating system levels and fibre channel adapter microcode levels. For this reason a detailed list of supported platforms is not included in this book. The current information is maintained on the Web site:

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

This Web site must always be consulted to ensure that only current information is used in any proposal or design.

The sales representative and/or customer is responsible for ensuring that the specific host system configuration used (that is, server model, operating system level and host adapter combination) is a valid and supported configuration by the server manufacturer.

Appendix B. One customer's situation

This report, written by Dave Reeve of IBM UK, is based on work carried out at the IBM Storage Center of Competence in Mainz to help a customer migrate from a SDG Fibre Channel implementation to native Fibre Channel attachment.

B.1 Introduction

The UK customer has a 2105-E20 attached to seven Netfinity 8000 systems. The Netfinity's have Qlogic 2100F Fibre Channel Host Bus Adapters (HBAs) connected to a pair of IBM 2108-S16 switches. The switches are attached to 3 x 2108-GO7 San Data Gateways which are then attached to eight ESS SCSI ports. The ESS runs a production 200GB SAP database. Figure 68 shows the customer environment.

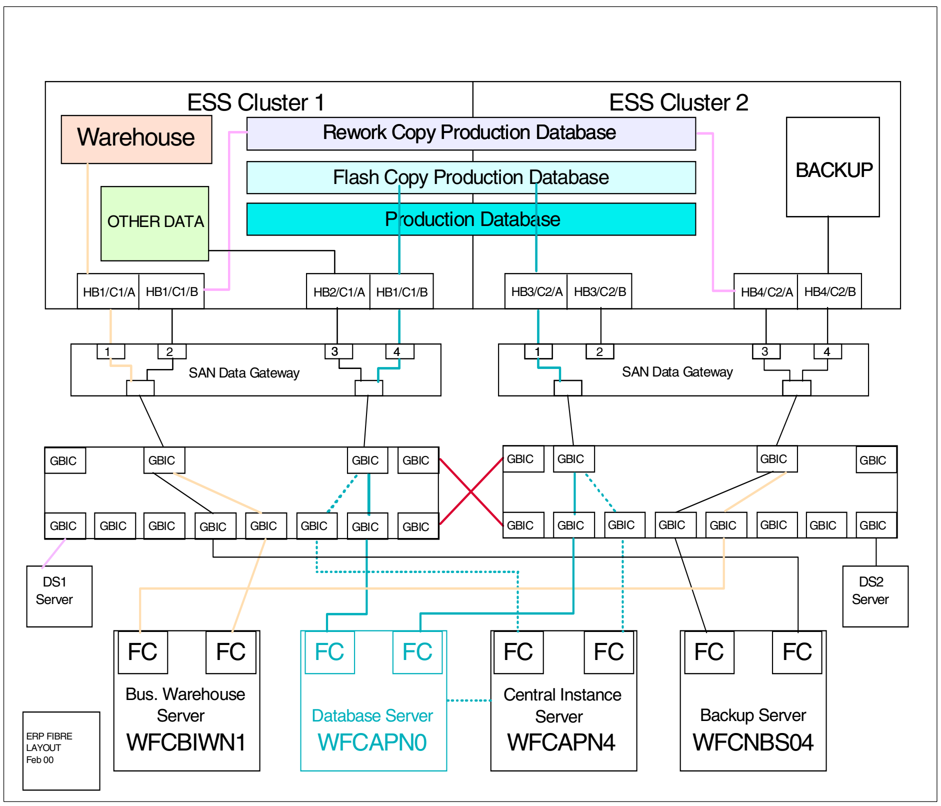

Figure 68. Customer environment

The customer had this environment in productive use and needed to understand how the migration from SDG to native adapters would be achieved and what disruption, if any, there would be to their applications. For volumes that were in production use, the overriding objective was to provide a seamless migration path from SDG to Native FC with fallback possible at all stages if things went wrong.

A migration test that was carried out at the Mainz SAN lab in June 2000. The migration was specifically for Windows NT, but most of the points will apply to UNIX systems as well. It discusses the principal steps involved in migration from SDG to FC adapters and provides details of the actual work carried out for this customer.

B.2 Test objectives

When the SDG is used, it maps the ESS logical volume SCSI IDs and targets to FCAL LUN IDs and presents these up the FC connection to the Netfinity starting at Lun_ID 1. Windows NT Disk Administrator then maps these Lun_IDs into disk numbers, which can be manually assigned to drive letters.

The migration from SDG to native FC attachment will cause the FC Lun_IDs to change. The customer wanted to ensure that the Windows NT drive letters would not change because of this, so the following test scenario was developed and performed at the IBM Mainz SAN Lab.

- Attach SDG to Netfinity using QL 2100F adapter
- Create twelve logical test volumes in the ESS
- Attach twelve test volumes to three ESS SCSI ports (four volumes per port)
- Scan Busses with SDG to create LUN_IDs
- Boot test NT system to get drive letters assigned to LUN_IDs
- Move twelve test volumes to ESS Fibre Channel port to create new LUN_IDs
- Attach Netfinity directly to ESS Fibre Channel port
- Boot test NT system and check that drive letters have remained the same even though LUN_IDs have changed

B.3 Migration planning

The migration exercise was split into a number of steps as summarized here. Each of these steps is explained in detail following this list:

- Fibre Adapter Considerations
- Install Fibre Adapters in the ESS
- Load new ESS microcode
- Migrate ESS volumes from SCSI port to Fibre port
- Connect ESS fibre ports to host HBAs
- Define ESS access profiles
- Reboot NT to see the volumes

B.3.1 Fibre adapter considerations

The ESS fibre channel adapter occupies one card slot in an I/O bay. If the ESS bay is fully populated with SCSI cards then it will be necessary to remove one of the SCSI cards first, so that the FC card can be installed. When this has to be done, it is very important that all logical volume assignments to both ports of the SCSI card are removed before it is replaced, otherwise unpredictable errors will occur.

If the volumes on the SCSI port are part of a multi-path host set and DPO or SDD is being used, then they can be dynamically removed from that adapter port with the ESS Specialist logical volume assignments panel. The remaining SCSI paths and the DPO/SDD S/W will ensure that the host system still has access to its data.

If the volumes on the SCSI port only have a single ESS connection then either the host system will have to lose access to data temporarily or the volumes and host must be assigned to another SCSI port. If the volumes can be assigned to another SCSI port with the same host type, then it may be acceptable to have these hosts share the same SCSI bus temporarily — although this might cause volume sharing issues. If attachment to another SCSI port is not possible, then the volumes will have to be removed from the SCSI port and left unassigned. Their contents will not be affected by being unassigned and once the FC adapters have been installed the volumes can be assigned to a new FC port with the ESS Specialist.

If a single path host has or is going to have a FC connection, then it may be possible to attach it to an existing SDG, either on a spare SDG port or through a FC switch into the SDG. In fact, if a switch is going to be used in the new environment it would be advantageous to install and test it prior to installing the ESS FC adaptors anyway. The host's volumes could then be attached to one of the SDG SCSI ports and the SDG VPS (Virtual Private SAN) software used to setup temporary LUN masking so that hosts do not see each other's data. Figure 69 shows the transition to this setup.

If there are spare slots in the ESS I/O bays then the FC card can be installed alongside the SCSI adapters. This has the benefit of making the migration steps easier and this is what was actually done for the customer migration test in the lab.

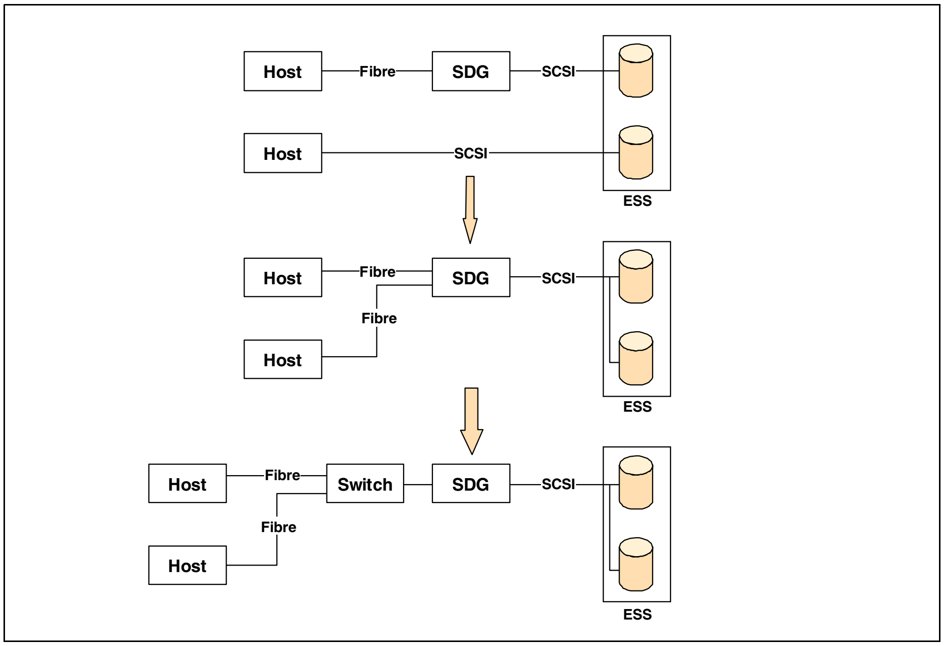

Figure 69. Migration steps

B.3.2 Install fibre adapters in the ESS

FC adapter card installation is a CE service action as the I/O bay has to be opened to allow host adapter cards to be removed and replaced. This means that any hosts with connections in that bay will be temporarily disconnected. If the host has DPO or SDD enabled as well as connections to other ESS host adapters, then there should be no loss of service. Hosts with single connections will need to be stopped before the I/O bay is taken out of service. The elapsed time taken from start to finish for the hardware install of the adapter card was approximately one hour per bay for the customer, which gives you an idea of how long a single host connection is likely to be interrupted.

B.3.3 Load new ESS microcode

Once the CE has installed the FC adapters the ESS microcode is updated. For most ESSs this will be a concurrent code load so no downtime is necessary. Some ESSs may not have the right level of code installed to allow a concurrent update — the CE should be contacted to understand what level the machine has. If downtime needs to be scheduled then it is also a good idea to install the FC adapters at the same time, because it will be much quicker than doing it concurrently.

B.3.4 Migrate ESS volumes to FC ports

There are two scenarios to consider concerning volume migration:

1. FC adapters can be installed alongside existing SCSI adapters.

   When the ESS FC adapters have been installed the logical volume assignment panels of the ESS Specialist can be used to modify the existing SCSI volume assignments. For the customer, the SCSI volumes were attached to the FC ports as well as the SCSI ports. This gave them the ability to retain the existing SDG setup which they knew worked and to test out the native FC connection in a non-pressured environment.

2. SCSI adapters must be removed to allow installation of FC adapters.

   The action taken here will depend upon whether or not volumes could be re-assigned to other SCSI ports. In the situation where temporary alternative host paths have been set up, then it is also possible to define volumes to the FC ports while they are connected to SCSI ports. If volumes had to be left unassigned then they will need to be assigned only to the FC ports.

   Note that when a volume is assigned to a FC port it becomes accessible to all the FC ports in the ESS. This is similar to a 3390 volume which is accessible through any ESCON port. In order to prevent undesired host access to volumes it is then necessary to setup an access profile for every host.

B.3.5 Setting up host to volume access profiles

The access profile for a volume is controlled by defining which HBAs are allowed to access it. Every HBA has a unique, hard coded ID known as the World Wide Port Name (WWPN), which has to be manually entered into the ESS Specialist host panel. This is different to the SDG, which can have an automatic host registration service provided by the SAN Data Gateway Explorer (which runs on the host systems).

If a new HBA is being installed on the host, then the WWPN ID can usually be obtained from the product documentation. However, if you are attaching to an existing FC HBA then you may need to discover what the HBA WWPN is. One method is to use the SDG Storwatch Specialist to look at the SAN connections prior to the FC install. If the SDG Explorer is running on an attached host system, then the WWPN of the HBA will be registered with the SDG and the WWPN can simply be read off the screen and entered into the ESS Specialist.

We did this for the customer but also used an alternative method which may be required for some installations. We booted the Netfinity and interrupted the boot at the time the Qlogic BIOS was being loaded. This presented us with a series of menu driven screens which enable various parameters to be changed. One of the screens shows the WWPN of the card installed and we verified that the ID shown here was the same as that presented via the San Data Gateway Explorer.

Once the WWPN is known it can be entered in the ESS Specialist host panel and the host added to the ESS host list via the *perform configuration update* function. In the test we only had one path to the Netfinity so we moved onto the next section. It is worth noting that if a host has multiple paths to the ESS then it will be necessary to define each HBA WWPN separately in these panels. A host with four HBAs would therefore have four host entries. This affects the number of hosts that can be attached to the ESS since there is a maximum of 512 WWPNs. For example, if every host had four HBAs then a maximum of 128 hosts could be attached.

B.3.6 Connect ESS fibre ports to host HBAs

The Netfinity servers are connected to 2109-S16 switches. There were spare GBICs on the switches which meant that is was possible to connect the ESS FC ports to the switch as well as the SAN Data Gateways. Soft zoning had already been set up in the switches between the NT system ports and the SDG ports, which meant that the NT systems had no access at this point in time to the ESS native FC ports.

New zones were then created in the switches between the NT systems and the ESS ports which provided a neat way of changing from a SDG connection to a native fibre connection. By just changing the active zone in the switch it was possible to flip to a native ESS FC configuration and then back to the SDG configuration if necessary. This met a key migration objective which was to have fallback possible at all stages of the migration.

If you don't have a switch but have the SDG connected directly to the host, then the same scenario can be created by running the fibre cable from the ESS FC port directly to the host FC port. The fibre connections can then be interchanged at the host port, although you will need to ensure that the ESS port is configured in loop mode since that is how the SDG operates.

B.3.7 Boot Windows NT to see new volumes

When the ESS was connected via the SDG's, we had used NT Disk Administrator to modify the Disk Volume Label to include the NT drive letter, for example, Customer_D0_L, where L is the drive letter. This was done so that we could easily check whether Drive L remained Drive L after the changeover.

The switch zoning was now changed to attach the Netfinity's to the ESS FC ports instead of the SDG's. Windows NT was rebooted and after the reboot had completed, we looked in Disk Administrator to check if the LUN_ID to drive letter mapping had stayed the same. We found that Drive L had a volume label of Customer_D0_L as before, and that all drives had been correctly mapped to their volume labels. This proved that although the disk LUN-IDs had changed because of the SDG to native adapter migration, it was completely transparent to Windows NT.

For the sake of our own interest, we then changed the switch zoning back to the SDG configuration and rebooted Windows NT. The drives were correctly mapped again as we had expected.

Appendix C. ESS Licensed Internal Code

The internal code name "GA+4" Licensed Internal Code (LIC) is being rolled out to the field and will be installed into all E20 model machines. For most customers there will be concurrent activation from the existing code. This must be checked with the IBM service engineer during the planning stage.

The code stream known as “1.1” is shipped with all F model machines. GA+4 and "1.1" are separate microcode streams but provide equivalent availability and subsystem operation enhancements, but 1.1 is required for advanced functions. The internal name "1.1" refers to the base LIC that provides XRC and platform for future advanced functions (FlashCopy, PPRC, native Fibre)

To install code level “1.1” onto an ESS that is at “GA+4” requires an interruption of service. However once the ESS is at 1.1 level, activation of future advanced functions can be done concurrently.

Appendix D. Special notices

This publication is intended to help IBM developers, Business Partners and customers who are involved with storage subsystems to specify, install and use ESS fibre attachment in UNIX and Windows NT environments. The information in this publication is not intended as the specification of any programming interfaces that are provided with the ESS. See the PUBLICATIONS section of the IBM Programming Announcement for ESS specifications and 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:

AIX	Netfinity
AS/400	OS/390
AT	RACF
CT	RAMAC
CUA	Redbooks
Current	Redbooks Logo
DFSMS/MVS	RS/6000
DFSMSdss	S/390
Enterprise Storage Server	Seascape
ESCON	SP
FICON	StorWatch
IBM ®	System/370
Lotus	System/390
Magstar	Wave
MVS/DFP	400
MVS/ESA	TME

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 E. Related publications

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

E.1 IBM Redbooks

For information on ordering these publications see , “How to get IBM Redbooks” on page 117.

- *IBM Enterprise Storage Server,* SG24-5465
- *Implementing the Enterprise Storage Server in Your Environment*, SG24-5420
- *Implementing ESS Copy Services on System/390*, SG24-5680
- *Introduction to Storage Area Network, SAN*, SG24-5470
- *IBM Storage Solutions for Server Consolidation,* SG24-5355
- *Storage Area Networks: Tape Future In Fabrics*, SG24-5474
- *Designing an IBM Storage Area Network,* SG24-5758
- *Planning and Implementing an IBM SAN*, SG24-6116
- *IBM Enterprise Storage Server Performance Monitoring and Tuning Guide*, SG24-5656
- *Using LCCM Functions with Servers and Workstations*, SG24-5292

E.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	SK2T-8043
IBM Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

E.3 Other resources

These publications are also relevant as a further information sources:

- *Designing Storage Area Networks: A Practical Reference for Implementing Fibre Channel SANs (The Addison-Wesley Networking Basics Series)*, ISBN 0201615843
- *IBM Enterprise Storage Server Web Users Interface Guide*, SC26-7346
- *IBM Storage Area Network Dat Gateway Installation and User's Guide*, SC26-7304

E.3.1 Referenced Web sites

These Web sites are also relevant as further information sources:

- `www.storage.ibm.com/ibmsan/index.htm` (IBM Enterprise SAN)
- `www.storage.ibm.com/hardsoft/products/fchub/fchub.htm` (IBM Fibre Channel Storage Hub)
- `www.pc.ibm.com/ww/netfinity/san` (IBM Storage Area Networks: NefinityServers)
- `www.storage.ibm.com/hardsoft/products/fcswitch/fcswitch.htm` (IBM SAN Fibre Channel Switch)
- `www.storage.ibm.com/hardsoft/products/fcswitch/supserver.htm` (IBM SAN Fibre Channel Switch)
- `www.ibm.com/storage/SANGateway` (IBM SAN Data Gateway)
- `www.storage.ibm.com/hardsoft/products/sangateway/supserver.htm` (IBM SAN Data Gateway)
- `www.storage.ibm.com/hardsoft/products/tape/ro3superserver.htm` (IBM SAN Data Gateway Router)
- `www.storage.ibm.com/hardsoft/products/fcss/fcss.htm` (IBMFibre Channel RAID Storage Server)
- `www.storage.ibm.com/hardsoft/products/ess/ess.htm` (Enterprise Storage Server)
- `www.storage.ibm.com/hardsoft/products/ess/supserver.htm` (Enterprise Storage Server)
- `www.brocade.com` (Brocade Communications Systems, Inc.)
- `www.fibrechannel.com` (Fibre Channel Industry Association)
- `www.mcdata.com` (McData Corporation)

- `www.pathlight.com` (Pathlight)
- `www.sanergy.com` (Tivoli SANergy)
- `www.snia.org` (Storage Networking Industry Association)
- `www.tivoli.com` (Tivoli)
- `www.t11.org` (Technical Committee T11)
- `www.vicom.com` (Vicom Systems)
- `www.vixel.com` (Vixel)
- `www.as.ibm.com/asus/connectivity.html` (IBM Global Services - connectivity)

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

This glossary contains a list of terms used within this redbook.

A

allegiance. The ESA/390 term for a relationship that is created between a device and one or more channel paths during the processing of certain condition.

allocated storage. On the ESS, this is the space that you have allocated to volumes, but not yet assigned.

application system. A system made up of one or more host systems that perform the main set of functions for an establishment. This is the system that updates the primary DASD volumes that are being copied by a copy services function.

AOM. Asynchronous operations manager.

APAR. Authorized program analysis report.

array. An arrangement of related disk drive modules that you have assigned to a group.

assigned storage. On the ESS, this is the space that you have allocated to volumes, and assigned to a port.

asynchronous operation. A type of operation in which the remote copy XRC function copies updates to the secondary volume of an XRC pair at some time after the primary volume is updated. Contrast with synchronous operation.

ATTIME. A keyword for requesting deletion or suspension at a specific target time.

availability. The degree to which a system or resource is capable of performing its normal function.

B

bay. Physical space on an ESS rack. A bay contains SCSI, ESCON or Fibre Channel interface cards and SSA device interface cards.

backup. The process of creating a copy of data to ensure against accidental loss.

C

cache. A random access electronic storage in selected storage controls used to retain frequently used data for faster access by the channel.

cache fast write. A form of fast write where the subsystem writes the data directly to cache, where it is available for later destaging.

CCA. Channel connection address.

CCW. Channel command word.

CEC. Central electronics complex.

channel. (1) A path along which signals can be sent; for example, data channel and output channel. (2) A functional unit, controlled by the processor, that handles the transfer of data between processor storage and local peripheral equipment.

channel connection address (CCA). The input/output (I/O) address that uniquely identifies an I/O device to the channel during an I/O operation.

channel interface. The circuitry in a storage control that attaches storage paths to a host channel.

channel path. The ESA/390 term for the interconnection between a channel and its associated controllers.

channel subsystem. The ESA/390 term for the part of host computer that manages I/O communication between the program and any attached controllers.

CKD. Count key data. An ES/390 architecture term for a device that specifies the format of and access mechanism for the logical data units on the device. The logical data unit is a track that can contain one or more records, each consisting of a count field, a key field (optional), and a data field (optional).

CLIST. TSO command list.

cluster. See storage cluster.

cluster processor complex (CPC). The unit within a cluster that provides the management function for the storage server. It consists of cluster processors, cluster memory, and related logic.

concurrent copy. A copy services function that produces a backup copy and allows concurrent access to data during the copy.

concurrent maintenance. The ability to service a unit while it is operational.

consistency group time. The time, expressed as a primary application system time-of-day (TOD) value, to which XRC secondary volumes have been updated. This term was previously referred to as “consistency time”.

consistent copy. A copy of data entity (for example a logical volume) that contains the contents of the entire data entity from a single instant in time.

contingent allegiance. ESA/390 term for a relationship that is created in a controller between a device and a channel path when unit-check status is accepted by the channel. The allegiance causes the controller to guarantee access; the controller does not present the busy status to the device. This enables the controller to retrieve sense data that is associated with the unit-check status, on the channel path with which the allegiance is associated.

control unit address (CUA). The high order bits of the storage control address, used to identify the storage control to the host system.

Note: The control unit address bits are set to zeros for ESCON attachments.

CUA. Control unit address.

D

daisy chain. A method of device interconnection for determining interrupt priority by connecting the interrupt sources serially.

DA. Device adapter.

DASD. Direct access storage device. See disk drive module.

data availability. The degree to which data is available when needed. For better data availability when you attach multiple hosts that share the same data storage, configure the data paths so that data transfer rates are balanced among the hosts.

data sharing. The ability of homogenous or divergent host systems to concurrently utilize information that they store on one or more storage devices. The storage facility allows configured storage to be accessible to any attached host systems, or to all. To use this capability, you need to design the host program to support data that it is sharing.

DDM. Disk drive module

data compression. A technique or algorithm that you use to encode data such that you can store the encoded result in less space than the original data. This algorithm allows you to recover the original data from the encoded result through a reverse technique or reverse algorithm.

data field. The third (optional) field of a CKD record. You determine the field length by the data length that is specified in the count field. The data field contains data that the program writes.

data record.A subsystem stores data records on a track by following the track-descriptor record. The subsystem numbers the data records consecutively, starting with 1. A track can store a maximum of 255 data records. Each data record consists of a count field, a key field (optional), and a data field (optional).

DASD-Fast Write. A function of a storage controller that allows caching of active write data without exposure of data loss by journaling of the active write data in NVS.

DASD subsystem. A DASD storage control and its attached direct access storage devices.

data in transit. The update data on application system DASD volumes that is being sent to the recovery system for writing to DASD volumes on the recovery system.

data mover. See system data mover.

dedicated storage. Storage within a storage facility that is configured such that a single host system has exclusive access to the storage.

demote. The action of removing a logical data unit from cache memory. A subsystem demotes a data unit in order to make room for other logical data units in the cache. It could also demote a data unit because the logical data unit is not valid. A subsystem must destage logical data units with active write units before they are demoted.

destage. (1) The process of reading data from cache. (2) The action of storing a logical data unit in cache memory with active write data to the storage device. As a result, the logical data unit changes from cached active write data to cached read data.

device. The ESA/390 term for a disk drive.

device address. The ESA/390 term for the field of an ESCON device-level frame that selects a specific device on a control-unit image. The one or two leftmost digits are the address of the channel to which the device is attached. The two rightmost digits represent the unit address.

device adapter. A physical sub unit of a storage controller that provides the ability to attach to one or more interfaces used to communicate with the associated storage devices.

device ID. An 8-bit identifier that uniquely identifies a physical I/O device.

device interface card. A physical sub unit of a storage cluster that provides the communication with the attached DDMs.

device number. ESA/390 term for a four-hexadecimal-character identifier, for example 13A0, that you associate with a device to facilitate communication between the program and the host operator. The device number that you associate with a subchannel.

device sparing. Refers to when a subsystem automatically copies data from a failing DDM to a spare DDM. The subsystem maintains data access during the process.

Device Support Facilities program (ICKDSF). A program used to initialize DASD at installation and perform media maintenance.

DFDSS. Data Facility Data Set Services.

DFSMSdss. A functional component of DFSMS/MVS used to copy, dump, move, and restore data sets and volumes.

director. See storage director and ESCON Director.

disaster recovery. Recovery after a disaster, such as a fire, that destroys or otherwise disables a system. Disaster recovery techniques typically involve restoring data to a second (recovery) system, then using the recovery system in place of the destroyed or disabled application system. See also recovery, backup, and recovery system.

disk drive module. The primary nonvolatile storage medium that you use for any host data that is stored within a subsystem. Number and type of storage devices within a storage facility may vary.

drawer. A unit that contains multiple DDMs, and provides power, cooling, and related interconnection logic to make the DDMs accessible to attached host systems.

DRAIN. A keyword for requesting deletion or suspension when all existing record updates from the storage control cache have been cleared.

drawer. A unit that contains multiple DDMs, and provides power, cooling, and related interconnection logic to make the DDMs accessible to attached host systems.

dump. A capture of valuable storage information at the time of an error.

dual copy. A high availability function made possible by the nonvolatile storage in cached IBM storage controls. Dual copy maintains two functionally identical copies of designated DASD volumes in the logical storage subsystem, and automatically updates both copies every time a write operation is issued to the dual copy logical volume.

duplex pair. A volume comprised of two physical devices within the same or different storage subsystems that are defined as a pair by a dual copy, PPRC, or XRC operation, and are in neither suspended nor pending state. The operation records the same data onto each volume.

E

ECSA. Extended common service area.

EMIF. ESCON Multiple Image Facility. An ESA/390 function that allows LPARs to share an ESCON channel path by providing each LPAR with its own channel-subsystem image.

environmental data. Data that the storage control must report to the host; the data can be service information message (SIM) sense data, logging mode sense data, an error condition that prevents completion of an asynchronous operation, or a statistical counter overflow. The storage control reports the appropriate condition as unit check status to the host during a channel initiated selection. Sense byte 2, bit 3 (environmental data present) is set to 1.

Environmental Record Editing and Printing (EREP) program. The program that formats and prepares reports from the data contained in the error recording data set (ERDS).

EREP. Environmental Record Editing and Printing

Program.

ERP. Error recovery procedure.

ESCD. ESCON Director.

ESCM. ESCON Manager.

ESCON. Enterprise Systems Connection Architecture. An ESA/390 computer peripheral interface. The I/O interface utilizes ESA/390 logical protocols over a serial interface that configures attached units to a communication fabric.

ESCON Director (ESCD). A device that provides connectivity capability and control for attaching any two ESCON links to each other.

extended remote copy (XRC). A hardware- and software-based remote copy service option that provides an asynchronous volume copy across storage subsystems for disaster recovery, device migration, and workload migration.

ESCON Manager (ESCM). A licensed program that provides host control and intersystem communication capability for ESCON Director connectivity operations.

F

F_Node Fabric Node - a fabric attached node.

F_Port Fabric Port - a port used to attach a NodePort (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.**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.

FBA. Fixed block address. An architecture for logical devices that specifies the format of and access mechanisms for the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size); the subsystem can access them independently.

FC-AL. Fibre Channel - Arbitrated Loop. An implementation of the fibre channel standard that uses a ring topology for the communication fabric.

FCS. See fibre channel standard.

fibre channel standard. An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. The protocol has two layers. The IP layer defines basic interconnection protocols. The upper layer supports one or more logical protocols (for example FCP for SCSI command protocols, SBCON for ESA/390 command protocols).

fibre channel ports. There are five basic kinds of ports defined in the Fibre Channel architecture, as well as some vendor-specific variations. The five basic ports are as follows:

Node Ports, N_ports. These ports are found in Fibre Channel Nodes, which are defined to be the source or destination of Information Units (IUs). I/O devices and host systems interconnected in point-to-point or switched topologies use N_ports for their connections. N_ports can only attach to other N_ports or to F_ports. The ESS Fibre Channel adapters support the N_port functionality when connected directly to a host or to a fabric.

Node-Loop Ports, NL_ports.These ports are just like the N_ports described above, except that they connect to a Fibre Channel Arbitrated Loop (FC-AL) topology. NL_ports can only attach to other NL_ports or to FL_ports. The ESS Fibre Channel adapters support the NL_port functionality when connected directly to a loop.

Fabric Ports, F_ports. These ports are found in Fibre Channel Switched Fabrics. They are not the source or destination of IUs, but instead function only as a "middleman" to relay the IUs from the sender to the receiver. F_ports can only attach to N_ports. The ESS Fibre Channel adapters do not support the F_port functionality, which is found only in fabrics.

Fabric-Loop Ports, FL_ports. These ports are just like the F_ports described above, except that they connect to an FC-AL topology. FL_ports can only attach to NL_ports. The ESS Fibre Channel adapters do not support the FL_port functionality, which is found only in fabrics or hubs.

Expansion Ports, E_ports. These ports are found in Fibre Channel Switched Fabrics and are used to interconnect the individual switch or routing elements. They are not the source or destination of IUs, but instead function like the F_ports and FL_ports to relay the IUs from one switch or routing element to another. E_ports can only attach to other E_ports. The ESS Fibre Channel adapters do not support the E_port functionality, which is found only in fabrics or hubs.

Combination E/F Ports, G_Ports. These ports are sometimes found in Fibre Channel Switched Fabrics and are used either as E_Ports, when the link is connected to another switch, or as F_Ports, when the link is connected to an N_Port for a host or device. This port automatically determines what mode to run in after determining what it is connected to.

fiber optic cable. A fiber, or bundle of fibers, in a structure built to meet optic, mechanical, and environmental specifications.

fixed utility volume. A simplex volume assigned by the storage administrator to a logical storage subsystem to serve as working storage for XRC functions on that storage subsystem.

FlashCopy. A point-in-time copy services function that can quickly copy data from a source location to a target location.

floating utility volume. Any volume of a pool of simplex volumes assigned by the storage administrator to a logical storage subsystem to serve as dynamic storage for XRC functions on that storage subsystem

G

GB. Gigabyte.

gigabyte. 1 073 741 824 bytes.

group. A group consist of eight DDMs. Each DDM group is a raid array.

GTF. Generalized trace facility.

H

HA. Home address, host adapter.

hard drive. A storage medium within a storage server used to maintain information that the storage server requires.

HDA. Head and disk assembly. The portion of an HDD associated with the medium and the read/write head.

HDD. Head and disk drive.

home address. A nine-byte field at the beginning of a track that contains information that identifies the physical track and its association with a cylinder.

host adapter. A physical sub unit of a storage controller that provides the ability to attach to one or more host I/O interfaces.

I

ICKDSF. See Device Support Facilities program.

identifier (ID). A sequence of bits or characters that identifies a program, device, storage control, or system.

IML. Initial microcode load.

initial microcode load (IML). The act of loading microcode.

I/O device. An addressable input/output unit, such as a direct access storage device, magnetic tape device, or printer.

I/O interface. An interface that you define in order to allow a host to perform read and write operations with its associated peripheral devices.

implicit allegiance. ESA/390 term for a relationship that a controller creates between a device and a channel path, when the device accepts a read or write operation. The controller guarantees access to the channel program over the set of channel paths that it associates with the allegiance.

Internet Protocol (IP). A protocol used to route data from its source to its destination in an Internet environment.

invalidate. The action of removing a logical data unit from cache memory because it cannot support continued access to the logical data unit on the device. This removal may be the result of a failure within the storage controller or a storage device that is associated with the device.

IPL. Initial program load.

ITSO. International Technical Support Organization.

J

JCL. Job control language.

Job control language (JCL). A problem-oriented language used to identify the job or describe its requirements to an operating system.

journal. A checkpoint data set that contains work to be done. For XRC, the work to be done consists of all changed records from the primary volumes. Changed records are collected and formed into a “consistency group”, and then the group of updates is applied to the secondary volumes.

K

KB. Kilobyte.

key field. The second (optional) field of a CKD record. The key length is specified in the count field. The key length determines the field length. The program writes the data in the key field. The subsystem uses this data to identify or locate a given record.

keyword. A symptom that describes one aspect of a program failure.

kilobyte (KB). 1 024 bytes.

km. Kilometer.

L

LAN. See local area network.

least recently used. The algorithm used to identify and make available the cache space that contains the least-recently used data.

licensed internal code (LIC).

(1) Microcode that IBM does not sell as part of a machine, but licenses to the customer. LIC is implemented in a part of storage that is not addressable by user programs. Some IBM products use it to implement functions as an alternative to hard-wired circuitry.

(2) LIC is implemented in a part of storage that is not addressable by user programs. Some IBM products use it to implement functions as an alternative to hard-wired circuitry.

link address. On an ESCON interface, the portion of a source, or destination address in a frame that ESCON uses to route a frame through an ESCON director. ESCON associates the link address with a specific switch port that is on the ESCON director. Equivalently, it associates the link address with the channel-subsystem, or controller-link-level functions that are attached to the switch port.

link-level facility. ESCON term for the hardware and logical functions of a controller or channel subsystem that allows communication over an ESCON write interface and an ESCON read interface.

local area network (LAN). A computer network located on a user's premises within a limited geographical area.

logical address. On an ESCON interface, the portion of a source or destination address in a frame used to select a specific channel-subsystem or control-unit image.

logical data unit. A unit of storage which is accessible on a given device.

logical device. The functions of a logical subsystem with which the host communicates when performing I/O operations to a single addressable-unit over an I/O interface. The same device may be accessible over more than one I/O interface.

logical disk drive. See logical volume.

logical subsystem. The logical functions of a storage controller that allow one or more host I/O interfaces to access a set of devices. The controller aggregates the devices according to the addressing mechanisms of the associated I/O interfaces. One or more logical subsystems exist on a storage controller. In general, the controller associates a given set of devices with only one logical subsystem.

logical unit. The SCSI term for a logical disk drive.

logical unit number. The SCSI term for the field in an identifying message that is used to select a logical unit on a given target.

logical partition (LPAR). The ESA/390 term for a set of functions that create the programming environment that is defined by the ESA/390 architecture. ESA/390 architecture uses this term when more than one LPAR is established on a processor. An LPAR is conceptually similar to a virtual machine environment except that the LPAR is a function of the processor. Also the LPAR does not depend on an operating system to create the virtual machine environment.

logical volume. The storage medium associated with a logical disk drive. A logical volume typically resides on one or more storage devices. A logical volume is referred to on an AIX platform as an hdisk, an AIX term for storage space. A host system sees a logical volume as a physical volume.

LSS. See logical subsystem.

LUN. See logical unit number.

least-recently used (LRU). A policy for a caching algorithm which chooses to remove the item from cache which has the longest elapsed time since its last access.

M

MB. Megabyte.

megabyte (MB). 1 048 576 bytes.

metadata. Internal control information used by microcode. It is stored in reserved area within disk array. The usable capacity of the array take care of the metadata.

million instructions per second (MIPS). A general measure of computing performance and, by implication, the amount of work a larger computer can do. The term is used by IBM and other computer manufacturers. For large servers or mainframes, it is also a way to measure the cost of computing: the more MIPS delivered for the money, the better the value.

MTBF. Mean time between failures. A projection of the time that an individual unit remains functional. The time is based on averaging the performance, or projected performance, of a population of statistically independent units. The units operate under a set of conditions or assumptions.

Multiple Virtual Storage (MVS). One of a family of IBM operating systems for the System/370 or System/390 processor, such as MVS/ESA.

MVS. Multiple Virtual Storage.

N

nondisruptive. The attribute of an action or activity that does not result in the loss of any existing capability or resource, from the customer’s perspective.

nonvolatile storage (NVS). Random access electronic storage with a backup battery power source, used to retain data during a power failure. Nonvolatile storage, accessible from all cached IBM storage clusters, stores data during

DASD fast write, dual copy, and remote copy operations.

NVS. Nonvolatile storage.

O

open system. A system whose characteristics comply with standards made available throughout the industry, and therefore can be connected to other systems that comply with the same standards.

operating system. Software that controls the execution of programs. An operating system may provide services such as resource allocation, scheduling, input/output control, and data management.

orphan data. Data that occurs between the last, safe backup for a recovery system and the time when the application system experiences a disaster. This data is lost when either the application system becomes available for use or when the recovery system is used in place of the application system.

P

path group. The ESA/390 term for a set of channel paths that are defined to a controller as being associated with a single LPAR. The channel paths are in a group state and are on-line to the host.

path-group identifier. The ESA/390 term for the identifier that uniquely identifies a given LPAR. The path-group identifier is used in communication between the LPAR program and a device to associate the path-group identifier with one or more channel paths. This identifier defines these paths to the control unit as being associated with the same LPAR.

partitioned data set extended (PDSE). A system-managed, page-formatted data set on direct access storage.

P/DAS. PPRC dynamic address switching.

PDSE. Partitioned data set extended.

peer-to-peer remote copy (PPRC). A hardware based remote copy option that provides a synchronous volume copy across storage subsystems for disaster recovery, device migration, and workload migration.

pending. The initial state of a defined volume pair, before it becomes a duplex pair. During this state, the contents of the primary volume are copied to the secondary volume.

pinned data. Data that is held in a cached storage control, because of a permanent error condition, until it can be destaged to DASD or until it is explicitly discarded by a host command. Pinned data exists only when using fast write, dual copy, or remote copy functions.

port. (1) An access point for data entry or exit. (2) A receptacle on a device to which a cable for another device is attached.

PPRC. Peer-to-peer remote copy.

PPRC dynamic address switching (P/DAS). A software function that provides the ability to dynamically redirect all application I/O from one PPRC volume to another PPRC volume.

predictable write. A write operation that can cache without knowledge of the existing formatting on the medium. All writes on FBA DASD devices are predictable. On CKD DASD devices, a write is predictable if it does a format write for the first record on the track.

primary device. One device of a dual copy or remote copy volume pair. All channel commands to the copy logical volume are directed to the primary device. The data on the primary device is duplicated on the secondary device. See also secondary device.

PTF. Program temporary fix.

R

RACF. Resource access control facility.

rack. A unit that houses the components of a storage subsystem, such as controllers, disk drives, and power.

random access. A mode of accessing data on a medium in a manner that requires the storage device to access nonconsecutive storage locations on the medium.

read hit. When data requested by the read operation is in the cache.

read miss. When data requested by the read operation is not in the cache.

recovery. The process of rebuilding data after it has been damaged or destroyed. In the case of remote copy, this involves applying data from secondary volume copies.

recovery system. A system that is used in place of a primary application system that is no longer available for use. Data from the application system must be available for use on the recovery system. This is usually accomplished through backup and recovery techniques, or through various DASD copying techniques, such as remote copy.

remote copy. A storage-based disaster recovery and workload migration function that can copy data in real time to a remote location. Two options of remote copy are available. See peer-to-peer remote copy and extended remote copy.

reserved allegiance. ESA/390 term for a relationship that is created in a controller between a device and a channel path, when a Sense Reserve command is completed by the device. The allegiance causes the control unit to guarantee access (busy status is not presented) to the device. Access is over the set of channel paths that are associated with the allegiance; access is for one or more channel programs, until the allegiance ends.

restore. Synonym for recover.

resynchronization. A track image copy from the primary volume to the secondary volume of only the tracks which have changed since the volume was last in duplex mode.

RVA. RAMAC Virtual Array Storage Subsystem.

S

SAID. System adapter identification.

SAM. Sequential access method.

SCSI. Small Computer System Interface. An ANSI standard for a logical interface to computer peripherals and for a computer peripheral interface. The interface utilizes a SCSI logical protocol over an I/O interface that configures attached targets and initiators in a multi-drop bus topology.

SCSI ID. A unique identifier assigned to a SCSI device that is used in protocols on the SCSI interface to identify or select the device. The number of data bits on the SCSI bus determines the number of available SCSI IDs. A wide interface has 16 bits, with 16 possible IDs. A SCSI device is either an initiator or a target.

Seascape architecture. A storage system architecture developed by IBM for open system servers and S/390 host systems. It provides storage solutions that integrate software, storage management, and technology for disk, tape, and optical storage.

secondary device. One of the devices in a dual copy or remote copy logical volume pair that contains a duplicate of the data on the primary device. Unlike the primary device, the secondary device may only accept a limited subset of channel commands.

sequential access. A mode of accessing data on a medium in a manner that requires the storage device to access consecutive storage locations on the medium.

server. A type of host that provides certain services to other hosts that are referred to as clients.

service information message (SIM). A message, generated by a storage subsystem, that is the result of error event collection and analysis. A SIM indicates that some service action is required.

sidefile. A storage area used to maintain copies of tracks within a concurrent copy domain. A concurrent copy operation maintains a sidefile in storage control cache and another in processor storage.

SIM. Service information message.

simplex state. A volume is in the simplex state if it is not part of a dual copy or a remote copy volume pair. Ending a volume pair returns the two devices to the simplex state. In this case, there is no longer any capability for either automatic updates of the secondary device or for logging changes, as would be the case in a suspended state.

SMF. System Management Facilities.

SMS. Storage Management Subsystem.

SRM. System resources manager.

SnapShot copy. A point-in-time copy services function that can quickly copy data from a source location to a target location.

spare. A disk drive that is used to receive data from a device that has experienced a failure that requires disruptive service. A spare can be pre-designated to allow automatic dynamic sparing. Any data on a disk drive that you use as a spare is destroyed by the dynamic sparing copy process.

SSA. Serial Storage Architecture. An IBM standard for a computer peripheral interface. The interface uses a SCSI logical protocol over a serial interface that configures attached targets and initiators in a ring topology.

SSID. Subsystem identifier.

stacked status. An ESA/390 term used when the control unit is holding for the channel; the channel responded with the stack-status control the last time the control unit attempted to present the status.

stage. The process of reading data into cache from a disk drive module.

storage cluster. A power and service region that runs channel commands and controls the storage devices. Each storage cluster contains both channel and device interfaces. Storage clusters also perform the DASD control functions.

storage control. The component in a storage subsystem that handles interaction between processor channel and storage devices, runs channel commands, and controls storage devices.

STORAGE_CONTROL_DEFAULT. A specification used by several XRC commands and messages to refer to the timeout value specified in the maintenance panel of the associated storage control.

storage device. A physical unit which provides a mechanism to store data on a given medium such that it can be subsequently retrieved. Also see disk drive module.

storage director. In an IBM storage control, a logical entity consisting of one or more physical storage paths in the same storage cluster. See also storage path.

storage facility. (1) A physical unit which consists of a storage controller integrated with one or more storage devices to provide storage capability to a host computer. (2) A storage server and its attached storage devices.

Storage Management Subsystem (SMS). A component of MVS/DFP that is used to automate and centralize the management of storage by providing the storage administrator with control over data class, storage class, management class, storage group, aggregate group and automatic class selection routine definitions.

storage server. A unit that manages attached storage devices and provides access to the storage or storage related functions for one or more attached hosts.

storage path. The hardware within the IBM storage control that transfers data between the DASD and a channel. See also storage director.

storage subsystem. A storage control and its

attached storage devices.

string. A series of connected DASD units sharing the same A-unit (or head of string).

striping. A technique that distributes data in bit, byte, multibyte, record, or block increments across multiple disk drives.

subchannel. A logical function of a channel subsystem associated with the management of a single device.

subsystem. See DASD subsystem or storage subsystem.

subsystem identifier (SSID). A user-assigned number that identifies a DASD subsystem. This number is set by the service representative at the time of installation and is included in the vital product data.

suspended state. When only one of the devices in a dual copy or remote copy volume pair is being updated because of either a permanent error condition or an authorized user command. All writes to the remaining functional device are logged. This allows for automatic resynchronization of both volumes when the volume pair is reset to the active duplex state.

synchronization. An initial volume copy. This is a track image copy of each primary track on the volume to the secondary volume.

synchronous operation. A type of operation in which the remote copy PPRC function copies updates to the secondary volume of a PPRC pair at the same time that the primary volume is updated. Contrast with asynchronous operation.

system data mover. A system that interacts with storage controls that have attached XRC primary volumes. The system data mover copies updates made to the XRC primary volumes to a set of XRC-managed secondary volumes.

system-managed data set. A data set that has been assigned a storage class.

T

TCP/IP. Transmission Control Protocol/Internet Protocol.

TOD. Time of day.

Time Sharing Option (TSO). A System/370 operating system option that provides interactive time sharing from remote terminals.

timeout. The time in seconds that the storage control remains in a "long busy" condition before physical sessions are ended.

timestamp. The affixed value of the system time-of-day clock at a common point of reference for all write I/O operations directed to active XRC primary volumes. The UTC format is yyyy.ddd hh:mm:ss.thmiju.

track. A unit of storage on a CKD device that can be formatted to contain a number of data records. Also see home address, track-descriptor record, and data record.

track-descriptor record. A special record on a track that follows the home address. The control program uses it to maintain certain information about the track. The record has a count field with a key length of zero, a data length of 8, and a record number of 0. This record is sometimes referred to as R0.

TSO. Time Sharing Option.

U

Ultra-SCSI. An enhanced small computer system interface.

unit address. The ESA/390 term for the address associated with a device on a given controller. On ESCON interfaces, the unit address is the same as the device address. On OEMI interfaces, the unit address specifies a controller and device pair on the interface.

Universal Time, Coordinated. Replaces Greenwich Mean Time (GMT) as a global time reference. The format is yyyy.ddd hh:mm:ss.thmiju.

utility volume. A volume that is available to be used by the extended remote copy function to perform data mover I/O for a primary site storage control's XRC-related data.

UTC. Universal Time, Coordinated.

V

vital product data (VPD). Nonvolatile data that is stored in various locations in the DASD subsystem. It includes configuration data, machine serial number, and machine features.

volume. An ESA/390 term for the information recorded on a single unit of recording medium. Indirectly, it can refer to the unit of recording medium itself. On a non-removable medium storage device, the terms may also refer, indirectly, to the storage device that you associate with the volume. When you store multiple volumes on a single storage medium transparently to the program, you may refer to the volumes as logical volumes.

vital product data (VPD). Information that uniquely defines the system, hardware, software, and microcode elements of a processing system.

VSAM. Virtual storage access method.

VTOC. Volume table of contents.

W

workload migration. The process of moving an application's data from one set of DASD to another for the purpose of balancing performance needs, moving to new hardware, or temporarily relocating data.

write hit. A write operation where the data requested is in the cache.

write miss. A write operation where the data requested is not in the cache.

write penalty. The term that describes the classical RAID write operation performance impact.

write update. A write operation that updates a direct access volume.

X

XDF. Extended distance feature (of ESCON).

XRC. Extended remote copy.

XRC planned-outage-capable. A storage subsystem with an LIC level that supports a software bitmap but not a hardware bitmap.

Index

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**	SG24-6113-00 Implementing Fibre Channel Attachment on the ESS
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/

Redbooks

Implementing Fibre Channel Attachment on the ESS

(0.2"spine)
0.17"<->0.473"
90<->249 pages

Implementing Fibre Channel Attachment on the ESS

Fast connectivity for a fast disk subsystem

Steps to migrate from SCSI connectivity to Fibre Channel

Positions the ESS at the heart of your SAN

The new Fibre Channel (FC) attachment on the ESS addresses connectivity issues. The limitations seen on SCSI in terms of distance, addressability and performance are gone with Fibre Channel attachment. The very rich connectivity options made possible by Fibre Channel technology have fueled the concept of SANs in the last year or two. Fibre Channel is the infrastructure of the SAN.

Fibre Channel attachment is not just a replacement of the parallel SCSI interface based on copper cables by a serial, fiber optic based interface. In view of today's operating system islands, it is more likely the entry into future universal data access.

This IBM Redbook will help you install, tailor and configure Fibre Channel (FC) attachment of Open Systems hosts to the IBM Enterprise Storage Server (ESS).

It gives a broad understanding of the procedures involved and describes the prerequisites and requirements, and then shows you how to implement FC attachment for each of the announced Open Systems platforms.

This redbook also describes the steps required to migrate to direct FC attachment from native SCSI adapters and from FC attachment via the SAN Data Gateway (SDG).

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

SG24-6113-00 ISBN 073841817X