IBM

Implementing IBM eserver xSeries SANs

- Set up a switched fabric connecting to an IBM FAStT500 Storage Server
- Configure FAStT Storage Manager in a Windows 2000 environment
- Support Windows and Linux servers on a single SAN

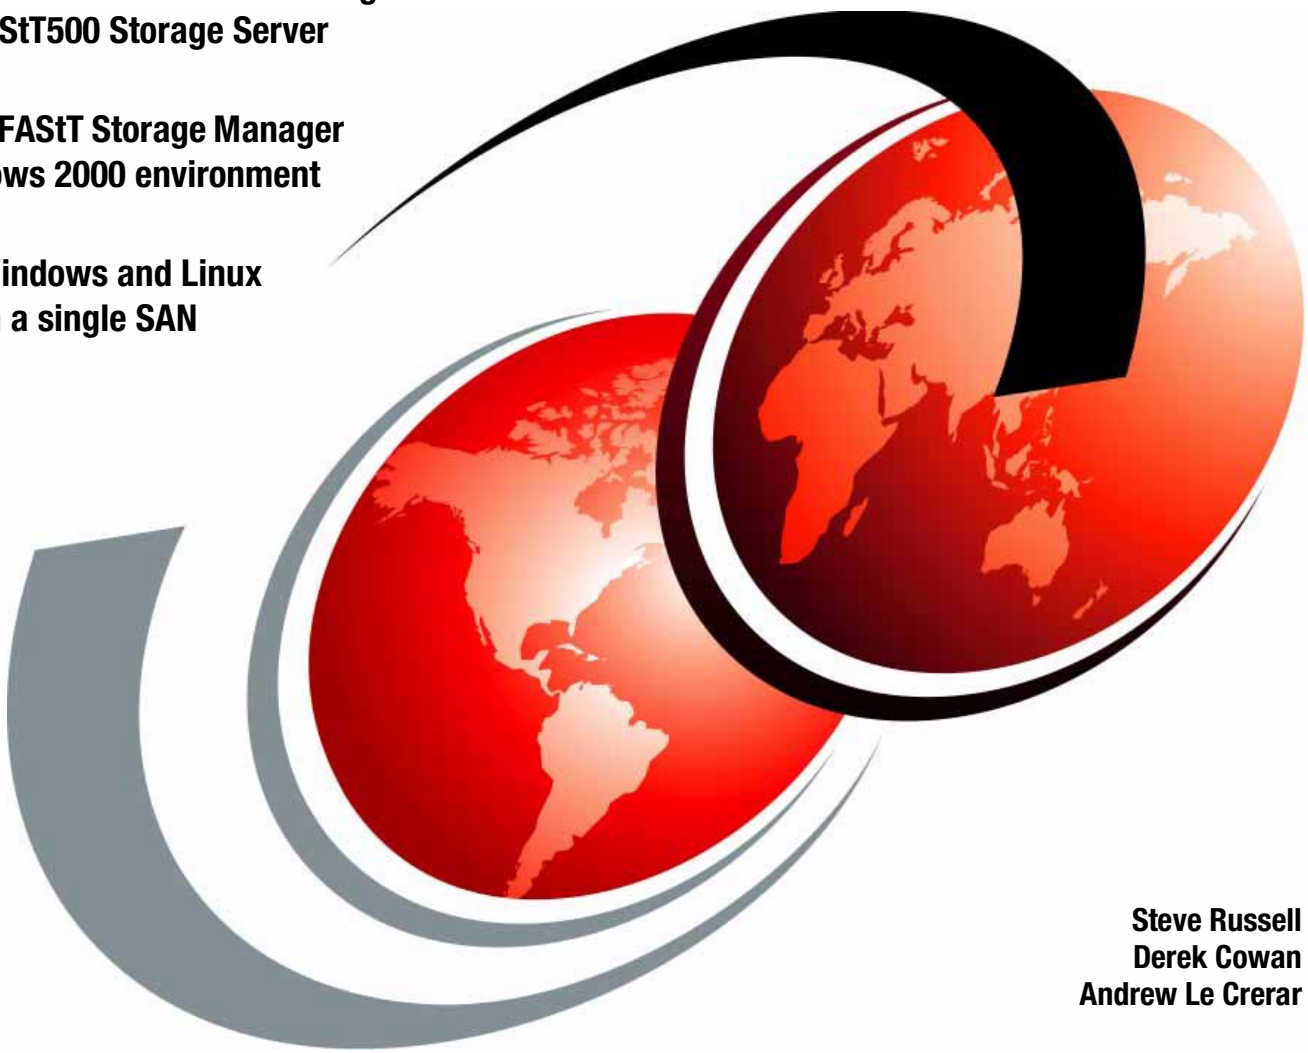

Steve Russell
Derek Cowan
Andrew Le Crerar

Redpaper

ibm.com/redbooks

International Technical Support Organization

Implementing IBM @server xSeries SANs

January 2002

Take Note! Before using this information and the product it supports, be sure to read the general information in “Special notices” on page 123.

First Edition (January 2002)

This edition applies to IBM Fibre Channel products designed for use with IBM @server xSeries servers.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 662
P.O. Box 12195
Research Triangle Park, NC 27709-2195

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

© Copyright International Business Machines Corporation 2002. All rights reserved.
Note to U.S Government Users - Documentation related to restricted rights - Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

© Copyright IBM Corp. 2002

Preface

Today, more than ever before, there are a number of factors that are driving businesses to consider different ways of providing storage for their servers. The number of servers being used within a company, or even within a department, often grows rapidly as business processes are computerized. This can cause inefficiencies in the way storage capacity is distributed and a lack of flexibility, accompanied by escalating management costs. Storage area networks (SANs) offer a solution to the quandary of having to meet the storage needs of your business in a way that is flexible, cost effective, and relatively simple to manage.

By offering a centralized pool of storage that is accessible by a number of different servers, capable of being distributed according to the need of individual systems, and that can be scaled as the needs of your business grow, and by providing sophisticated management tools, storage area networks offer an attractive alternative to server-specific storage subsystems.

These benefits may easily offset costs that are, initially at least, higher than those for direct-attached storage. Some additional complexity in configuration also has to be considered, but the advantages already mentioned, and the levels of redundancy available, which ensure reliable operation, make a compelling argument for implementing SANs in any network having more than a handful of servers.

This Redpaper will help you install, tailor and configure a storage area network using IBM xSeries servers and the latest Fibre Array Storage Technology (FAStT) subsystems they support.

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

Steve Russell is a Certified Consulting IT Specialist at the International Technical Support Organization, Raleigh Center, where he manages residencies and produces Redbooks and Redpapers related to IBM xSeries servers. Before joining the ITSO, Steve worked in the Technical Marketing field in the UK as a member of the IBM Netfinity organization in EMEA. Prior to that, he spent nearly 15 years managing and developing PC-based hardware and software projects at the IBM Hursley laboratory in the UK. He holds a degree in Electrical and Electronic Engineering, and is a member of the Institution of Electrical Engineers and a Chartered Engineer.

Derek Cowan is a Technical Support Engineer for IBM Australia in the Asia Pacific Technical Support Center. He previously worked in the World-wide Level 2 Helpcentre in Greenock, Scotland. He has five years' experience in xSeries, Netfinity, and PC Servers. He holds an honours degree in Mechanical Electronic Systems from Glasgow Caledonian University. His areas of expertise include Windows NT 4.0, MSCS, SCSI, DASD and Storage Area Networking.

© Copyright IBM Corp. 2002

Andrew Le Crerar is a Systems Engineer working for Capespan (Pty) Ltd in Cape Town, South Africa. He has been involved with the FAStT product from its release and has attended storage training at IBM Greenock, Scotland. He is a Microsoft Certified Systems Engineer and an IBM @server xSeries Systems Expert. His areas of expertise include Windows NT 4.0, Windows 2000, MSCS, Terminal Services, Storage Area Networking, IBM xSeries and Netfinity.

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

Linda Robinson	IBM International Technical Support Organization, Raleigh Center
Paul Chenger	xSeries Briefing Center, IBM Raleigh
Ronnie Winick	IBM Executive Briefing Center
Farrel Benton	xSeries Server Product Engineering
Mic Watkins	xSeries FAStT Storage Products Development

Notice

This publication is intended to help IBM technical staff, business partners, and customers to implement storage area networks based on IBM FAStT technology. The information in this publication is not intended as the specification of any programming interfaces that are provided by xSeries servers or FAStT disk subsystems. See the PUBLICATIONS section of the relevant IBM Programming Announcements for these products for more information about what publications are considered to be product documentation.

IBM trademarks

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

e (logo)® @
DB2®
FlashCopy™
IBM®
Netfinity®
Perform™
Redbooks™
Redbooks Logo
ServeRAID™
SP™
SP2®
Tivoli®
xSeries™
z/OS™
Lotus®

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this Redpaper or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  `ibm.com/redbooks`

- Send your comments in an Internet note to:

  `redbook@us.ibm.com`

- Mail your comments to the address on page ii.

1

Storage area network overview

In today's business environment, the growth in e-business and data-intensive applications such as data warehousing and data mining is pushing the limits of disk subsystem performance and capacity. Fibre Channel is an architecture designed to meet these increasing demands.

As often found with new technology, customers sometimes display uncertainty and a reluctance to deploy unfamiliar products. Since one of the main strengths of Fibre Channel is its flexibility, it allows you to create complex configurations that can be confusing if you are new to the subject. This Redpaper is intended to overcome some of the complexity issues and offer guidance to assist you in installing and configuring an IBM SAN solution using xSeries servers. Along the way, we demonstrate how Fibre Channel architecture is meeting the new demands already mentioned.

In this opening chapter, we describe a SAN and the benefits it brings, and discuss some of the terminology associated with SANs.

1.1 SAN Introduction

Most current Intel CPU-based servers are directly attached to their own storage subsystem. Typically, the CPU communicates with attached disks using a Small Computer System Interface (SCSI) bus. In contrast, a storage area network (SAN) is a topology that interconnects servers with a separate storage subsystem. The components of a SAN are switches, hubs, gateways and routers. These are familiar names to anyone who works with local area networks (LANs), and a storage area network can be thought of as a LAN for the disk subsystem. As this is a relatively new network architecture, some interoperability issues can arise when connecting multiple vendors' devices into a single SAN. IBM offers a single vendor solution, which eliminates these concerns.

© Copyright IBM Corp. 2002

To help you to understand exactly what comprises a SAN, let us split this new technology into three main components:

- SAN fabric

  Fibre Channel defines a physical structure for communication between, for example, servers and a disk subsystem and is often the technology of choice when constructing a SAN. The interconnecting network of devices such as switches, hubs, gateways and routers is referred to as the SAN fabric.

- SAN servers

  Being able to provide centralized storage for a number of servers is the primary reason that SAN solutions are attractive, since this offers flexibility, cost effectiveness and ease of management.

  Systems connected in a SAN can include a mix of server platforms, such as Windows 2000, UNIX (of various flavors) and z/OS. As customers drive towards server consolidation and e-business, the need for SAN solutions will increase. Although the early SAN solutions supported only homogeneous (single operating system) environments, SANs have now evolved to the point where they can support a truly heterogeneous (multiple operating system) environment.

- SAN storage

  Storage is the resource on which your valuable information resides, and it must therefore support your company's business objectives. Simply deploying more and faster storage devices is not enough; a new kind of infrastructure is needed, one that provides enhanced network availability, data accessibility, and system manageability than that offered by traditional storage subsystems.

Figure 1-1 shows an example of the topology for a network that utilizes a SAN for server storage. In it, we see the network servers attached to disks and other storage subsystems over a storage area network:

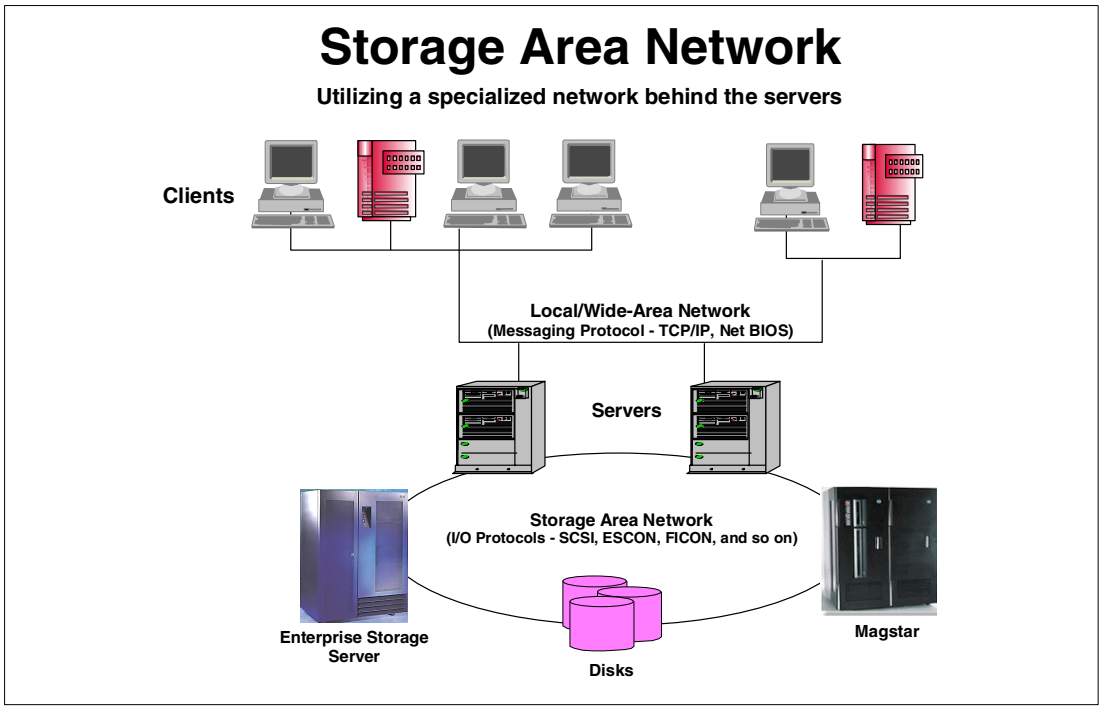

Figure 1-1 Topology of a network using SAN storage

1.2 Small Computer Systems Interface (SCSI)

Fibre Channel has close links with parallel SCSI as it supports the SCSI protocol. Therefore, we feel that it should be discussed. In doing so, we explain SCSI's limitations and give reasons as to why Fibre Channel is the new high-end disk architecture.

Even today SCSI is the predominant server disk interface technology used in xSeries servers. SCSI was designed for the personal and small computer environment, and became an American National Standards Institute (ANSI) standard in 1986. It allows attachment of peripheral devices such as disk drives, CD-ROMs and tapes, using a bus architecture that is most commonly 16 bits wide. SCSI can be split into two components, namely the physical interconnection among devices and the protocol or command language by which devices communicate.

Although very successful and popular, the SCSI architecture does have some limitations. These include support for a maximum of 16 devices on the bus, and a maximum cable length of 12 meters. SCSI bus speeds have increased significantly since its inception, but further increases to the bus clocking speed (to boost performance) is becoming more and more difficult, for a number of reasons. One problem, for example, is that higher bus clock rates generally mean shorter maximum bus lengths must be observed, due to the properties of signal propagation along the bus. Configuring large amounts of online disk storage while observing such restrictions can be challenging.

That said, SCSI is still the mainstream server disk technology and many xSeries systems include basic SCSI support as standard on the system board. For users who wish to benefit from the security of RAID technology, IBM offers the IBM ServeRAID-4 family of controllers.

1.3 Fibre Channel

As the demands on storage have increased, the limitations of parallel SCSI have become more pressing, and a new architecture became necessary in order to overcome them. Fibre Channel addresses these issues in a number of ways. Being a serial architecture means that some of the signal propagation problems with handling high-speed parallel data do not arise. In addition, using optical transmission techniques rather than electrical signals means that cable lengths can be significantly longer. (Fibre Channel can be implemented using copper cables but, even then, supports longer cables than SCSI.)

Fibre Channel is an ANSI standard that uses a serial transmission rate of 1 GHz and an aggregate data rate of 200 MBps (both soon to be doubled), compared to SCSI's maximum of 80 MHz and 160 MBps. The SCSI command protocol is well defined and is used over Fibre Channel for controlling disks, although other protocols are also supported over a Fibre Channel link.

Three distinct topologies are supported by Fibre Channel:

- Point-to-point
- Arbitrated loop
- Switched fabric

These are shown in Figure 1-2:

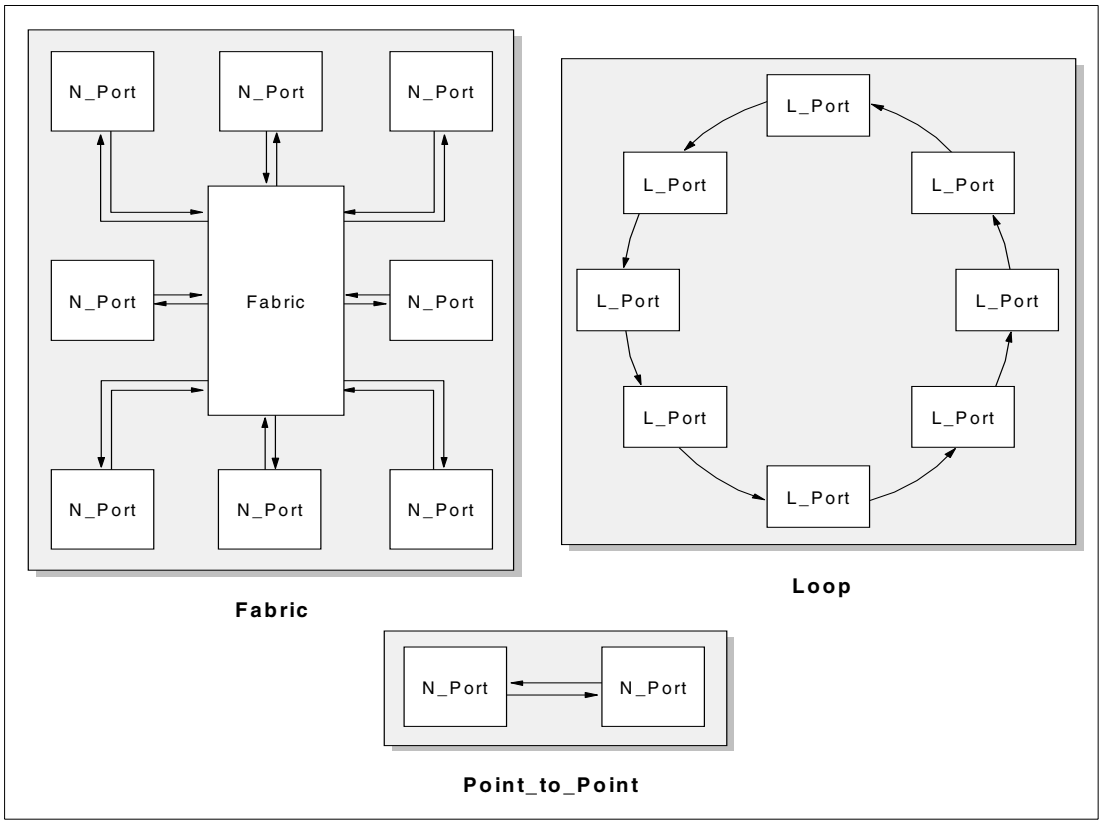

Figure 1-2 Fibre Channel topologies

Let us take a look at each of them, in turn.

- Point-to-point

  This is the simplest topology and is a straightforward connection between two Fibre Channel nodes. The two devices utilize the total bandwidth of the connection.

- Arbitrated loop

  Fibre Channel Arbitrated Loop (FC-AL) is a loop topology that supports up to a maximum of 126 devices, all sharing the available bandwidth. Data is passed to a device, which examines the data and which, if it is intended for that device, keeps it; otherwise the data is passed onto the next device in the loop.

  Since this is an arbitrated topology, only two devices can talk on the loop at any one time. A device will request access to the loop and, if successful, a connection is made between the sender and a receiver. Frames are transferred around the loop using the SCSI protocol. Once the transfers have finished, the communication ends and the loop is again available for arbitration by other devices. To ensure that other devices can win arbitration, a fairness algorithm can be employed.

- Switched fabric

  Fibre Channel Switched (FC-SW) topology is typically used in SANs. Full bandwidth dedicated links are available at each port of a switch. Thus, when you add a device to the switch, you are actually increasing the aggregate bandwidth. This is in contrast to arbitrated loops, in which the existing bandwidth has to be shared among all attached devices, and thus adding a device effectively reduces the average bandwidth available to other devices on the loop. A switched fabric can currently support up to 16 million devices.

1.4 IBM xSeries servers and Fibre Channel

As Fibre Channel technology has evolved, IBM, as a technology leader, has embraced it and released product into the marketplace. We can broadly break the introduction of Fibre Channel products into three phases:

- Phase 1

  In 1998 IBM introduced the 3526, which was the first Fibre Channel RAID controller for IBM Intel-based servers. It supported a Fibre Channel link between server and controller, but the disks it drove were attached using a SCSI bus. The 3526 was accompanied by the 3523 unmanaged hub. Using the unmanaged hub and the RAID controller, you could create arbitrated loop topologies.

  The 3526 allowed customers to migrate some of their parallel SCSI disk enclosures into a Fibre Channel environment. The IBM 3526 had six SCSI parallel ports, allowing up to 60 drives to be attached, and host-side fiber connections that allowed you to connect to an unmanaged hub. This meant that you could attach your SCSI disks to a SAN's fabric. The maximum number of drives that could be connected to a single SCSI adapter at the time was 30, using the ServeRAID-3HB SCSI RAID Controller. The 3526 introduced the Intel-based server market to the advantages of Fibre Channel, namely support for increased attachment distance, greater bandwidth and higher total storage capacity.

- Phase 2

  The next step was to have an end-to-end fiber solution. In 2000, IBM released Fibre Array Storage Technology (FAStT) products such as the FAStT500 RAID controller, EXP500 expansion unit, and 2109-S08 and 2109-S16 Fibre Channel switches, so that this could now be achieved. With the ability to implement fiber connections all the way down to the individual disk drives, we now have end-to-end connectivity at 1 Gbps. Also, with the release of the 2109-S08 and -S16 switches, you may now implement a switched SAN fabric, with a higher aggregate bandwidth than that available in a hub-based network.

  The FAStT500 with its dual loop capability allows you to attach up to 22 storage enclosures for a maximum of 220 drives. With 73.4 GB drives supported, this means a total storage capacity of 16.1 TB. With the ability of the 2109 switches to cascade, this means that storage scalability is now a reality.

- Phase 3

  While this paper was being prepared, IBM released the FAStT700 Storage Server. The FAStT700 supports 2 Gbps host-side connections and uses small form factor gigabit interface converters (GBICs). The FAStT700 has two controllers each with a 700 MHz processor and 1 GB of SDRAM cache memory.

  The new FAStT FC-2 Host Bus Adapter is a 64-bit 66 MHz PCI-X adapter that supports 2 Gbps Fibre Channel connections. The adapter can automatically negotiate for Fibre Channel bit rates of 1 or 2 Gbps. The card is also compatible with 32-bit and 33 MHz PCI slots.

  The EXP700 is a new 14-bay Fibre Channel disk enclosure and supports 1 Gbps drives initially, with future support for 2 Gbps drives.

  In order to enable implementation of complete 2 Gbps fabrics, IBM will also be releasing a 2109 F16 switch where each port supports 2 Gbps connections.

  With the availability of these new products, IBM has also released Version 8 of the IBM Storage Manager software. This version has some new features, such as the ability to increase a logical drive's capacity and Flashcopy, which creates a logical point-in-time image of a physical disk and can require less disk space as it is a logical copy not a physical one. Typically this feature would be used for backup applications to read data and perform a backup while the original drive is still online for user access.

1.5 Advantages of SANs

Storage area networks have many benefits and advantages over traditional direct-attached storage technologies. The following list highlights the most important of these:

- **Increased bandwidth**

  A Fibre Channel connection typically supports a transfer rate of 100 MBps half-duplex and 200 MBps full-duplex. The latest Fibre Channel products run at 200 MBps half-duplex and 400 MBps full-duplex. The maximum rates supported by SCSI at the moment are supported by Ultra160 devices, with a 160 MBps transfer rate. 320 MBps products are planned for the future.

- **Scalability**

  Fibre Channel scales extremely well. With the use of switches, you can keep adding more and more devices to the SAN. Another factor in SAN scalability is support for much greater drive capacity. The ServeRAID-4 SCSI RAID controllers with fully populated EXP300s can have a maximum of 56 drives connected to a single controller. The FAStT500 can have a maximum of 220 drives attached to it.

- **Longer distances**

  SCSI Ultra160 connections support a maximum distance of 12 meters. Short-wave Fibre Channel can support distances up to 500 meters. If even greater distances between server and storage are needed, long-wave Fibre Channel can support distances of up to 10 km. By cascading switches, this distance can be increased up to 70 km. This distance is limited by the maximum number of hops (seven) allowed between devices.

- **Centralized storage and management**.

  SAN technology allows you to have a pool of storage from which you can allocate disk capacity to servers as required. This reduces the amount of unused capacity inherent in direct-attached storage configurations, where changing storage capacity requires adding or removing individual drives accessible by a single server. Extra storage can also be added without shutting the systems down. With SAN Navigator you have the ability to totally manage your SAN infrastructure.

- **Reduced LAN traffic**.

  Data backup is critical, but backing up each server with its own tape, or across the network with a shrinking backup window, is either costly or affects LAN performance. With a SAN, you can perform LAN-free and server-free backups with tape pooling.

- **High availability**

  SCSI cabling limits common disk subsystem clusters to only two machines. With a SAN providing storage, this limitation is removed. Using the FAStT500 product, we now have the ability to have clusters with four, eight, 16, and 32 machines as the FAStT500 currently supports up to 16 storage partitions. Switch cascading allows you to eliminate any single point of failure in your fabric, further enhancing high availability.

1.6 Terminology

As is typical when a new technology is introduced, a number of new and unfamiliar terms are used when discussing it. The following list defines those you are most likely to encounter:

- **LUN**

  Logical Unit Number. To an operating system, a LUN appears as a disk drive. A LUN can be a physical disk, a tape, or a logical disk when using RAID. The LUN is a unit of storage that is assigned to a *host*.

- **LUN masking**

  LUN masking is the ability to hide LUNS (drives) from other hosts. It is particularly useful in a *heterogeneous environment*.

- **Host**

  This is the term that represents a computer attached to the SAN.

- **Host port**

  A *host port* is a physical connection to the PCI adapter that allows an xSeries server a fiber connection into the SAN. The server can be attached, using its host port, to a switch or hub or directly to a FAStT500 control unit.

- **Host Group**

  A *host group* is a collection of hosts. The members of a host group are usually defined to have access to common storage (that is, they are in a clustered configuration).

- **Homogeneous environment**

  Your environment is said to be *homogenous* when only one operating system is in use by servers in the SAN.

- **Heterogeneous environment**

  Your environment is said to be *heterogeneous* when multiple operating systems are in use by different servers in the SAN, for example Windows 2000 and Red Hat Linux.

- **Zoning**

  Zoning allows segmentation of a SAN at the switch level. This segmentation can be created at either a software or hardware level.

  - Hardware zoning uses physical ports and groups these to form a zone.
  - Software zoning uses World Wide Names (*WWNs*) to form a zone.

- **Fabric**

  A *fabric* is the switched network that interconnects the Fibre Channel devices belonging to the SAN.

- **WWN**

  Every Fibre Channel device is given a *World Wide Name*, which is an IEEE-registered 64-bit unique number.

- **Cascading**

  Cascading is the term used to refer to interconnected switches.

- **GBIC**

  Gigabit Interface Converters are hot-pluggable modules connecting the electronic interfaces of Fibre Channel devices to the medium used in the SAN fabric. GBICs are available to support short-wave and long-wave optical fiber, and also copper cables.

2

Planning a SAN

Installing a SAN is a relatively complex undertaking that requires proper planning. In this chapter we discuss some the steps you should take before starting to construct your own SAN. During the writing of this paper, we implemented a small SAN in our lab and we give examples of our own planning steps in the hope that these will assist you with your own implementations.

2.1 Hardware

First of all we look at the hardware that we used for our SAN. The discussion is intentionally brief. If you require more details about the products to which we refer in this section, we highly recommend *Fibre Array Storage Technology, A FAStT Introduction*, SG24-6246.

2.1.1 FAStT500 Storage Server

The FAStT500 Storage Server, shown in Figure 2-1 on page 10, is a 4U-high, dual Fibre Channel RAID controller. Each controller has a 300 MHz AMD K6 processor with 256 MB of cache, which can be upgraded to 512 MB. The FAStT500 has fiber connectivity all the way from host to drives and supports 1 Gbps transfer rates. Each controller has a serial connection and a 10/100 Mbps Ethernet port for management.

The FAStT500 supports four host-side mini-hubs and another four mini-hubs for connections to drive enclosures. On the drive side, each mini-hub supports a single loop. Two mini-hubs provide a single FC-AL loop with redundancy, and the full complement of four mini-hubs can therefore support two redundant loops.

A single Fibre Channel loop can support 126 devices. However, since each EXP500 external disk enclosure houses 10 drives, all of which require an ID, and a management chip, which also uses an ID, a total of 11 enclosures, using 121 IDs, can be attached to a loop. The total number of drives supported by two loops is therefore 220.

© Copyright IBM Corp. 2002

Figure 2-1 FAStT500 Storage Server

2.1.2 FAStT EXP500 Storage Expansion Unit

The EXP500 is a 3U external racj-mountable storage enclosure which supports a maximum of 10 Fibre Channel disj drives. The rear of the enclosure has two Enclosure Services Monitor (ESM) boards. These boards have LEDs to indicate errors and events. The ESM boards also have Tray ID switches. These switches need to be set to identify the enclosure when multiple units are attached to a FAStT Storage Server.

The EXP500 has a SCSI Access Fault Tolerant Enclosure (SAF-TE) chip, which is a management device that is used for monitoring and management of both the enclosure and the drives it contains. A storage enclosure fault could be, for example, a fan failure or a high temperature warning. An EXP500 enclosure is shown in Figure 2-2:

Figure 2-2 EXP500 Storage Expansion Unit

2.1.3 SAN Fibre Channel switches

IBM offers two SAN Fibre Channel switches, the 2109 Models S16 and S08. Common features of these switches are that they both have hot-swappable redundant power supplies and cooling fans, and both have an RJ-45 connector supporting 10/100 Ethernet connections. This port allows remote telnet and Web access for configuration, monitoring and testing of the switch.

Both models also run the same fabric operating system, which allows the interconnection of switches to create a SAN with hundreds of ports, providing high performance, scalability and fault tolerance.

Differences between the switches are that the S16 is a 2U high device with 16 ports. The S16 also has an LED panel and a keypad, which are used to configure an IP address to allow you to use a Web browser and JAVA for management.

In contrast, the S08 is a 1U high device that has eight ports. A terminal emulation program, such as HyperTerminal, is used to perform IP address configuration over a serial port connection. This allows you to use a Web browser and JAVA for management.

Figure 2-3 shows a picture of the IBM 2109 S16 switch:

Figure 2-3 IBM 2109 S16 switch

2.1.4 FAStT Host Bus Adapter

The FAStT Host Bus Adapter is a 64-bit PCI adapter that can run at 66 MHz. The card is backward compatible and can work in a PCI slot that is 32 bit and 33 MHz speed. This adapter supports only short-wave Fibre Channel connections, since the GBIC is fixed.

The adapter is compliant with PCI Version 2.2 and Fibre Channel-Arbitrated Loop (FC-AL-2) standards, and supports point-to-point and fabric connections. It support Fibre Channel to SCSI and IP protocols, but only Fibre Channel to SCSI is supported in xSeries server solutions. The adapter is shown in Figure 2-4:

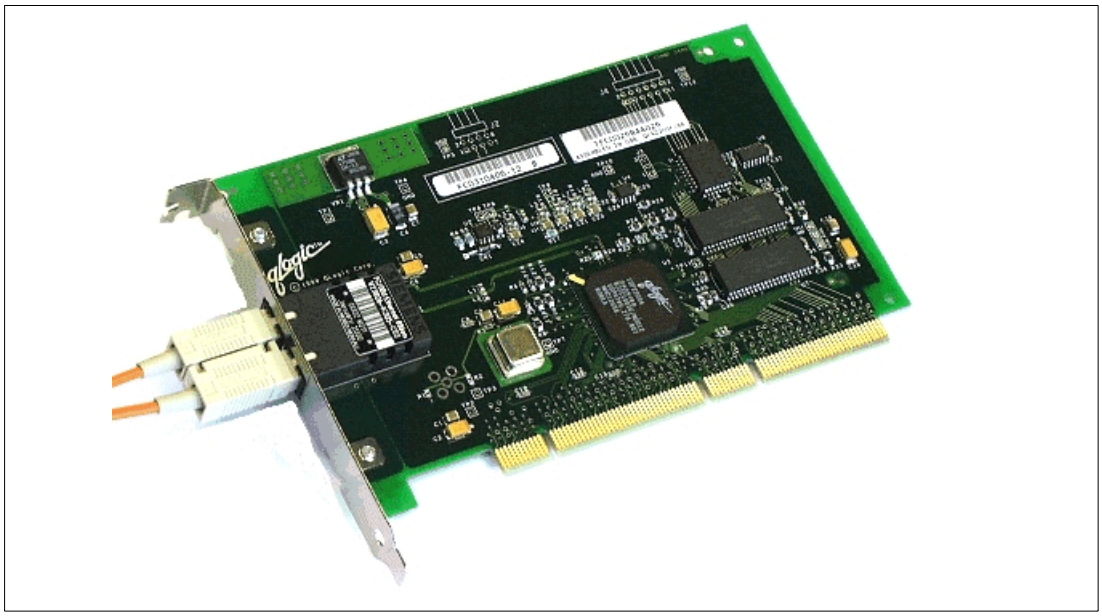

Figure 2-4 FAStT Host Bus Adapter

2.2 Configuring the hardware

Correct configuration is the single most important factor in implementing a SAN. This section breaks the process down into its major steps, introducing them in sequence. Later chapters then provide more detail about these steps.

2.2.1 Cabling hardware

Cabling the hardware is an extremely important part of the install process. You should define a naming standard and label all cables accordingly. We recommend reading Chapter 2 of the *Fibre Array Storage Technology, A FAStT Introduction*, SG24-6246, and Chapter 2 of *Netfinity FAStT500 RAID Controller Enclosure Unit Installation Guide*, shipped with the product, for detailed information about how to cable your FAStT500 and EXP500 units.

In our lab configuration, we started with a simple SAN, based on an xSeries server with two HBAs installed, a switch, a FAStT500, and one EXP500. A schematic for this configuration is shown in Figure 2-5:

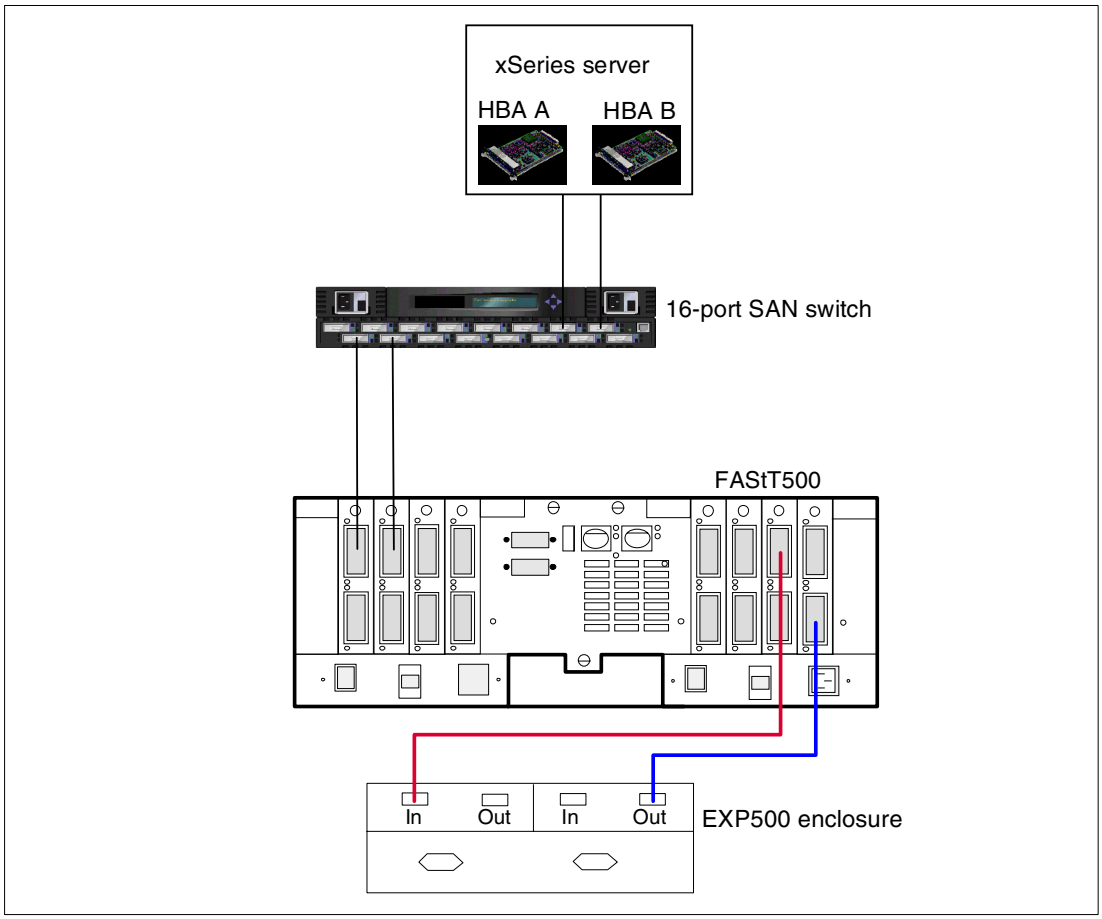

Figure 2-5 Basic SAN configuration

In our experience, a first point of confusion for people new to Fibre Channel hardware is understanding how connecting one cable can create a loop. To understand this we have to look more closely at the cable itself, as shown in Figure 2-6:

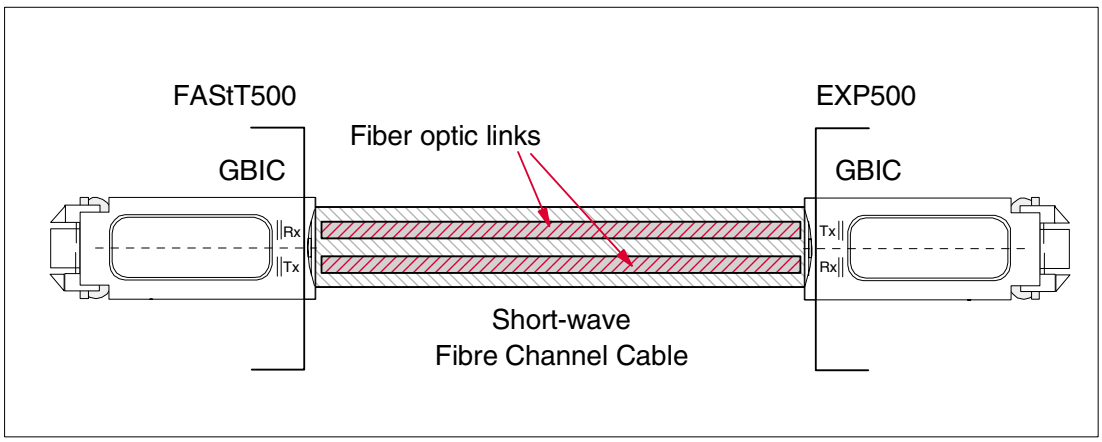

Figure 2-6 Cable diagram

The cable has two fiber connectors, a transmit and a receive connector for each device. The transmit port of one device is connected to the receive port of the other. Therefore, when, for example, you connect the cable from a RAID Controller to an EXP storage enclosure, a loop is formed.

To improve overall system availability, we recommend that you install two FAStT HBAs in each xSeries server. Doing so allows you to create redundant paths and helps to ensure that your data is always accessible. We believe that if you are going to make a significant investment in a high-performance solution, then the cost of a second adapter is more than offset by the improved availability you achieve.

2.2.2 Power on and off sequence

The following procedure must be followed when applying or removing power to an xSeries SAN solution. The FAStT500 controllers require that the attached EXP500 storage units are powered on before power is applied to the FAStT500 Storage Server. This is due to the controllers' configuration being stored on the drives themselves, and is specific to the Tray IDs and Slot IDs of the attached EXP500 storage units.

Important: If you do not power on in the correct sequence you could lose your storage server configuration and possibly your data.

The recommended power-on procedure is:

1. Apply power to your EXP500 storage units
2. Apply power to your Fibre Channel hubs or switches
3. Apply power to your FAStT500 Storage Server
4. Apply power to your servers

Important: If you do not power off in the correct sequence you could lose your storage server configuration and possibly your data.

The recommended power-off procedure is:

1. Remove power from your servers.
2. Remove power from your FAStT500 Storage Server
3. Remove power from your Fibre Channel hubs or switches
4. Remove power from your EXP500 storage units

2.2.3 Zoning

Zoning is a technique used in the SAN switches in order to segment your SAN. In a Windows environment, if you do not zone properly you may have multiple paths to the drives, causing a single LUN to appear multiple times under the operating system. This is to be avoided, so we zone the switch in order that each FAStT HBA has only one path to a RAID controller in the FAStT500.

When it is time to zone your IBM switches, you have to decide whether to implement soft zoning, hard zoning, or a combination of both. Let us examine these zoning methods:

- Hardware zoning

  Hardware zoning groups physical ports on a switch to form a zone. The advantage of hardware zoning is that, if a device fails, a replacement can be connected into the same port and no additional administration work is required. If you select this approach, adequate physical security measures must be taken to prevent unauthorized individuals from making connections to configured ports, and thereby accessing your SAN.

- Software zoning

  Software zoning uses unique World Wide Names (WWNs) to form your zone. This method of zoning allows you to attach a meaningful alias name to each device, making your zone easier to administer. As WWNs are used, you can physically change the ports to which devices are attached without disrupting the SAN configuration. A minor negative aspect to software zoning is that there is some administration overhead if an adapter fails and needs replacing. If and when this occurs, you have to reassign the old alias name to the new adapter.

In our configuration we chose software zoning. This enabled us to scale and implement redundancy easily during the switch upgrade we undertook later. When we cascaded the switches and moved connections to the new switch, we did not have to do anything to the zone configuration. The security and configuration integrity was maintained without any administration work.

Once you have correctly zoned your SAN, allowing each adapter only one path to a controller, the redundant disk array controller (RDAC) driver is used to provide the operating system with fault-tolerant paths to the disks. The RDAC essentially communicates with the RAID controllers, requesting them to move the disks across to the RAID controller with which it can communicate when the original path to the LUN fails.

For example, referring back to Figure 2-5 on page 13, if we assume that HBA A is zoned with controller A, HBA B is zoned with controller B, and that the D: drive on the xSeries server is owned by Controller A, what happens when Controller A now fails? The RDAC driver tries to read data from the D: drive and finds its path blocked. So the RDAC uses HBA B to communicate with Controller B and requests that the drive be transferred. Controller B then takes ownership of the drive and then the RDAC is able to fulfill the read request from the D: drive.

2.3 Pre-installation download of code

As we go through the installation process in the following chapters, a series of device drivers, firmware, BIOS upgrades, and utilities are required. As with all good planning you should review the firmware and drivers available so that everything is prepared in advance. To close this chapter, we list the code that you require and from where it can be obtained. We also document the code levels we used in our lab configuration.

2.3.1 FAStT Host Bus Adapter installation

The latest code versions are available from:

http://www.pc.ibm.com/support

At the time we wrote this paper, you could download the code by following these steps:

1. From the list of systems below the Browse heading, click **Servers** to go to the xSeries support home page.
2. In the Family listbox, select **Fibre Channel Solutions** to go to the Fibre Channel Solutions support page.
3. On the left-hand navigator, click **Downloadable files**.
4. In the Downloadable files by category listbox, select **Fibre Channel Solutions** and you will get a listing of all Fibre Channel files.

Click the relevant links from this list to download the files you require. The files we suggest you should check include:

- The device driver for the FAStT Host Bus Adapter. We downloaded Fibre Channel Solutions - IBM FAStT Host Adapter Device Driver for Microsoft Windows 2000 Version 8.00.09.06.
- The firmware for the FAStT Host Bus Adapter. We downloaded Servers - IBM FAStT Host Adapter BIOS Version 1.68.
- The IBM FAStT Check utility. We downloaded Fibre Channel Solutions - IBM FAStT Check Application Version 2.0.

The *IBM FAStT Host Adapter Installation and User's Guide* can be located at:

`http://www.pc.ibm.com/qtechinfo/MIGR-4WJHJX.html`

2.3.2 Switch installation

The preferred management tool for managing the switch is Web browser/JAVA-based rather than using a serial interface or telnet. Java 2 Version 1.3.1_01 can be obtained from:

`http://java.sun.com/j2se/1.3/download-windows.html.`

To ensure that you can perform all necessary switch functions you must have the latest switch firmware. This can be downloaded from:

`http://www.storage.ibm.com/ibmsan/products/2109/download.htm`

In our case we used level `a2.4.1c`. To perform the upgrade you need some additional software, which is contained in a file called rsh.zip, available from the same site.

2.3.3 FAStT500 and FAStT Storage Manager setup

The latest code versions are available from:

`http://www.pc.ibm.com/support`

You can locate them in the same way as indicated for the Host Bus Adapter files, described in 2.3.1, "FAStT Host Bus Adapter installation" on page 15. From the list of files, you should download:

- The IBM FAStT Fibre Channel Controller firmware. We downloaded Fibre Channel Solutions - IBM FAStT Fibre Channel Controller Firmware Version 04.01.02.07. This package file contains the Fibre Channel Controller firmware and NVSRAM.
- The IBM FAStT Storage Manager for Microsoft Windows NT and Windows 2000. We downloaded Fibre Channel Solutions - IBM FAStT Storage Manager for Microsoft Windows NT and Windows 2000 Version 7.10. This package file contains the Storage Manager Client, RDAC and Utilities.

2.3.4 Checklist

Table 2-1 lists the levels of code we used in our lab environment. You may find a similar table useful in your planning:

Table 2-1 Code levels used in our lab environment

Product	Type	Code Level
FAStT HBA	Firmware	1.68
FAStT HBA	Device Driver	8.00.09.06
FAStT Check	Utility	2.0
Switch	JAVA	1.3.1_01
Switch	Firmware	a2.4.1c
FAStT500	Firmware	04.01.02.07
FAStT500	NVSRAM	NV3552R710NT004
Storage Manager 7.10	Client	07.10.G5.05
Storage Manager 7.10	RDAC	07.10.95.03
Storage Manager 7.10	Utility	07.10.25.01

3

FAStT Host Bus Adapter installation

In this chapter we document the steps taken to complete the installation of the FAStT Host Bus Adapters in our xSeries server. For details of the physical installation procedure, please refer to the documentation provided with your IBM xSeries PCI Host Bus Adapter. The *IBM Installation and User's Guide* for the adapter can be located at:

http://www.pc.ibm.com/qtechinfo/MIGR-4WJHJX.html

Important:

If you have an IBM ServeRAID adapter installed in your xSeries server, you should ensure that the IBM xSeries FAStT Host Bus Adapter is installed in a higher PCI slot. The onboard SCSI or ServeRAID adapter should be seen during POST before the FAStT Host Bus Adapter.

Please refer to the relevant xSeries server *Hardware Maintenance Manual* to obtain PCI slot information for your specific xSeries server model.

3.1 IBM xSeries PCI Host Bus Adapter firmware

After physically installing the two adapters, you must ensure that they are at the latest level of firmware. Using the diskette created from the code downloaded earlier:

1. To identify the current firmware level, boot the xSeries server and press <ALT> <Q> when you see the adapter's BIOS messages during POST, as shown in Figure 3-1:

© Copyright IBM Corp. 2002

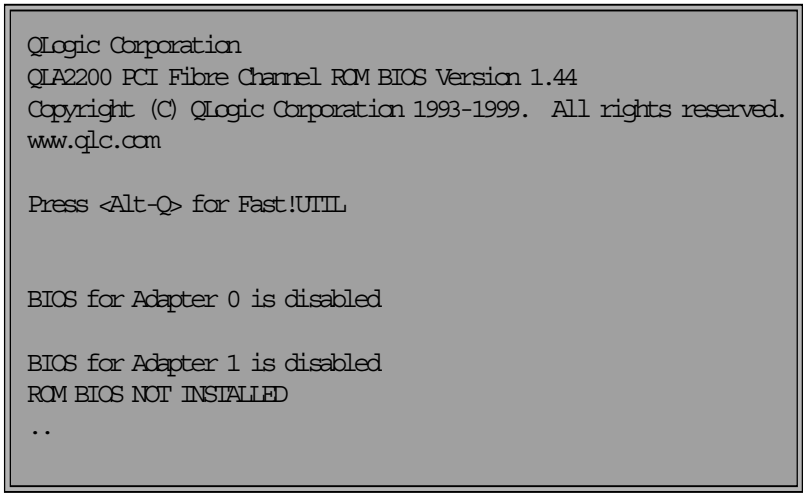

Figure 3-1 IBM xSeries PCI Host Bus Adapter POST

2. At the Main Menu screen, select **Configuration Settings** and then select **Host Adapter Settings**, to display the adapter configuration information as shown in Figure 3-2:

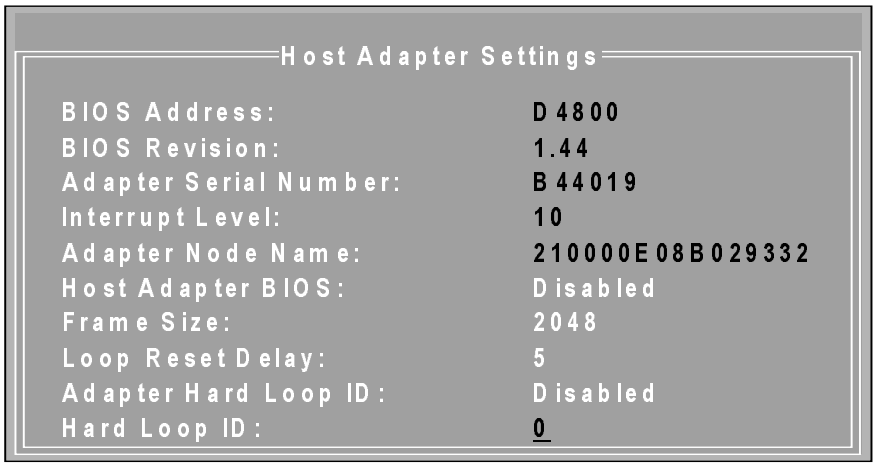

Figure 3-2 Host Adapter Settings

The second item listed is the BIOS Revision number.

3. It is always useful to have a check point where you can confirm that what you have done so far is working well. Now is a good time to ensure that all your hardware is connected correctly. If your machine is able to detect the IBM FAStT500, then this confirms that your adapter, cables, and switch are working correctly.

   From the Main Menu, select **Scan Fibre Devices**. If two entries, one for each controller, are displayed for the IBM 3552 FAStT500 unit, you know that the basic connections are good.

4. Insert the firmware update disk in the system's diskette drive, exit the BIOS Menu and the server will reboot.

Note: Before proceeding, ensure that you read the readme.txt file on the flash diskette to check for any last-minute changes or important information.

5. When the server boots you will be presented with a command prompt. At the command prompt, run **`flasutil /f /l`**.
6. Once this is complete you will again be shown a prompt. Enter **`flasutil /U`**.

> **Important:** Ensure that you use the /U switch. This will ensure that both adapters are updated and IBM defaults are set.

7. Now if you return to the Host Adapter Settings in the adapter's BIOS you should see the new BIOS version listed.

3.2 Host Bus Adapter device driver installation

With the adapters physically installed and at the latest firmware level, you now need to be make them available to the operating system by installing the appropriate device driver:

1. Boot to the operating system. In our case this was Microsoft Windows 2000 Advanced Server. The operating system should detect the cards automatically.
2. Select **My Computer->Manage->Device Manager->SCSI and RAID controllers.** The adapters are listed in Device Manager as shown in Figure 3-3:

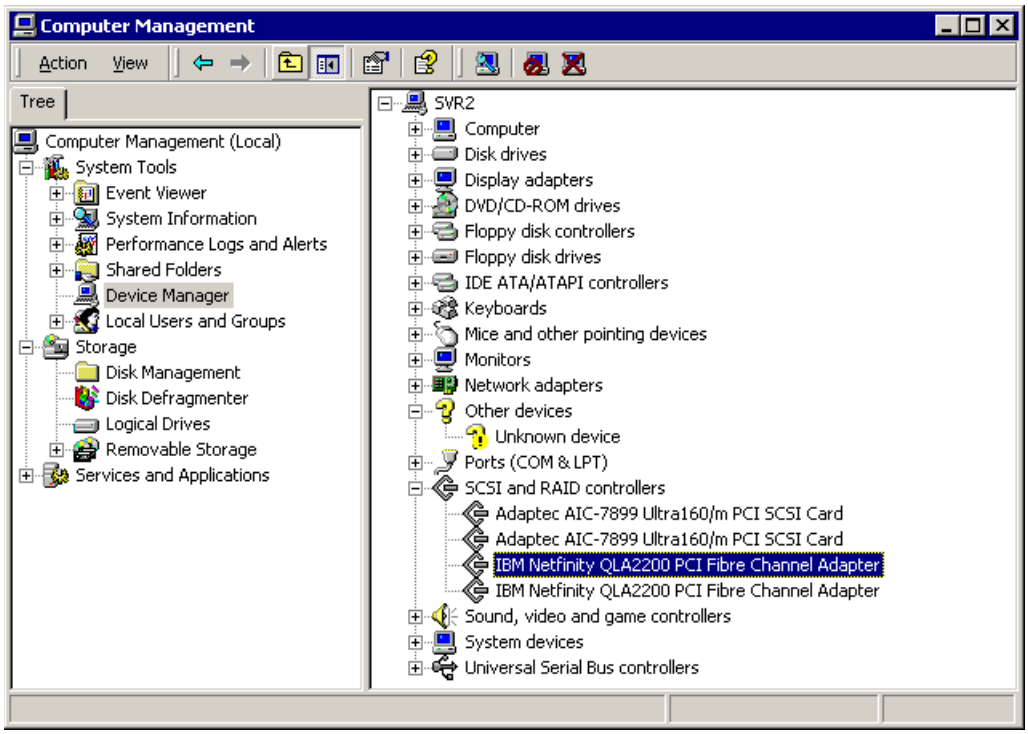

Figure 3-3 Windows automatic device detection

3. Highlight the device as shown in Figure 3-3, right-click and select **Properties**. Select the **Driver** tab and you can check the Driver Version as shown in Figure 3-4:

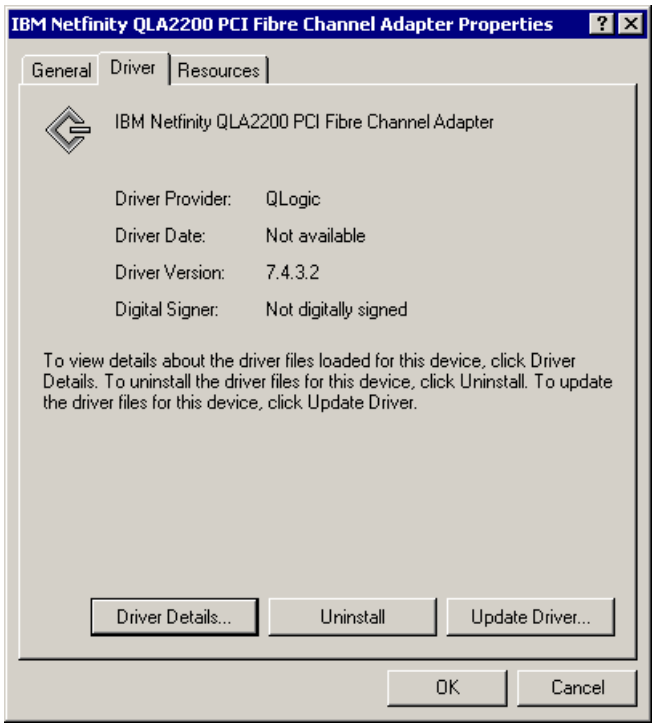

Figure 3-4 Properties window for the host adapter

4. Now click **Update Driver** and install the driver that you downloaded in the planning stage. Follow the steps, selecting **Have Disk** and enter the path to the new driver as shown in Figure 3-5:

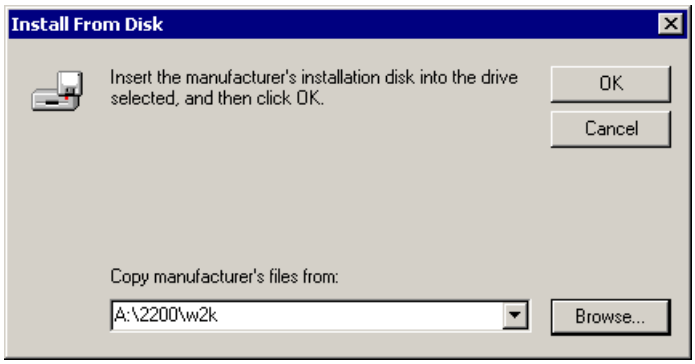

Figure 3-5 Pointing to the new driver

5. After the update is complete, the change is reflected in the Device Manager. For our example, as shown in Figure 3-6, the adapter detail has changed from IBM Netfinity QLA2200 PCI Fibre Channel Adapter to QLogic QLA2200 PCI Fibre Channel Adapter:

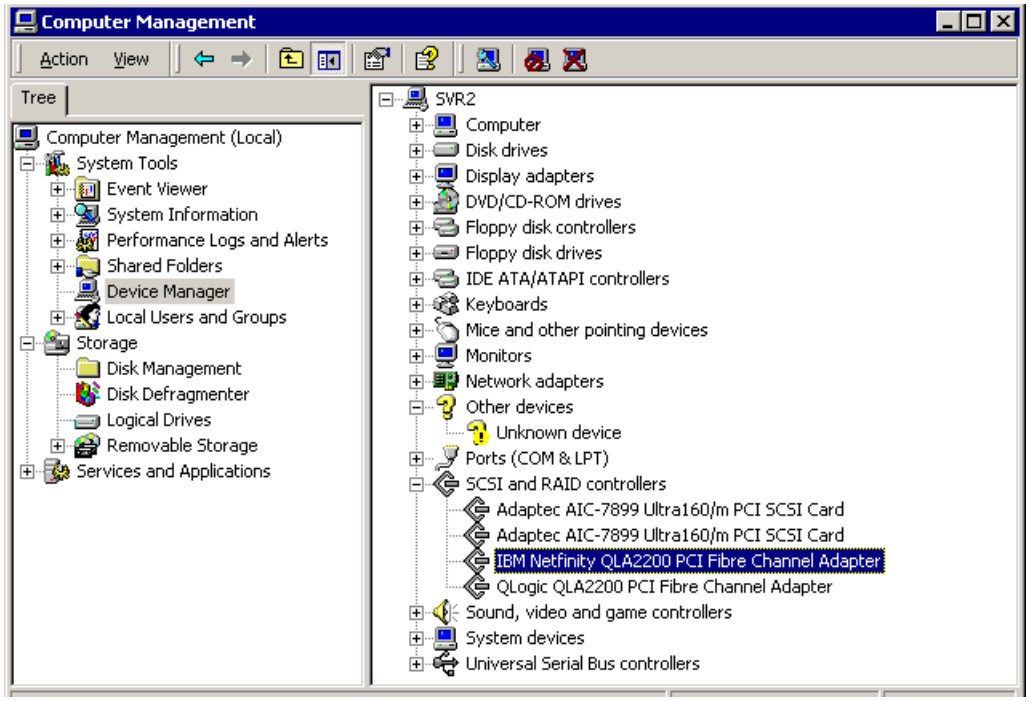

Figure 3-6 Device Manager after only one driver updated

6. Repeat steps 3 to 5 above for the second IBM xSeries FAStT Host Bus Adapter in the machine. Now Computer Management should appear as shown in Figure 3-7:

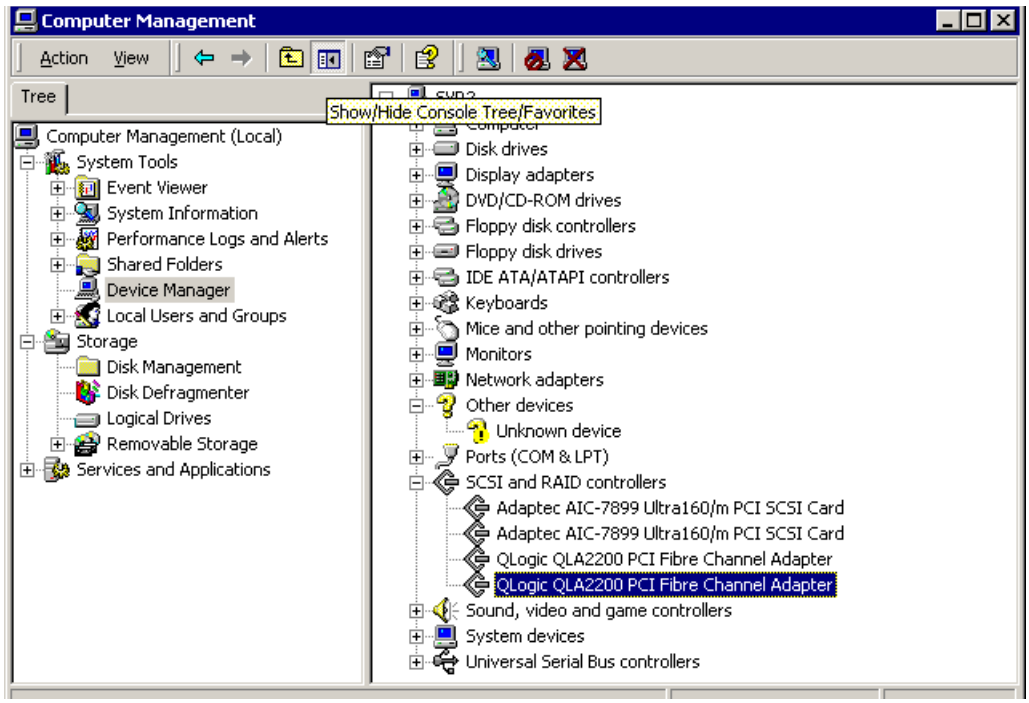

Figure 3-7 IBM xSeries FAStT Host Bus Adapter driver update completed

7. You can confirm that the correct driver level is now installed by checking the driver details as before. Figure 3-8 shows the new driver version after we installed it:

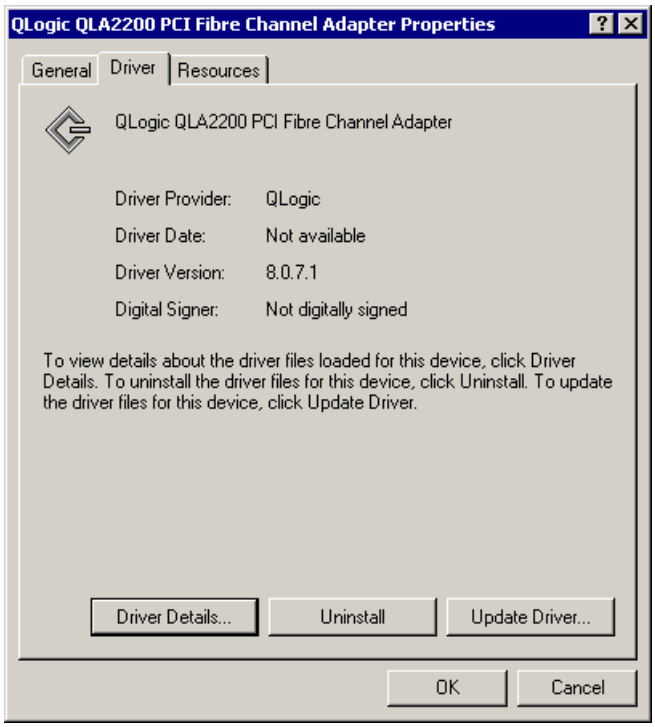

Figure 3-8 Updated driver details

Installation of IBM FAStT Host Bus Adapters is now complete.

3.3 IBM FAStT Check application

This utility is a valuable management tool. There are two components, an agent and a GUI. We installed both on our servers. The FAStT Check utility has the following features:

- It provides you with general information about the adapters, such as the PCI slot and bus number used by an adapter. It also gathers WWN information and can be used to identify ports and HBAs when creating zones or allocating storage partitions to hosts.
- A number of useful statistics are available:
  - Adapter Errors: The number of adapter errors reported by the adapter driver.
  - Device Errors: The number of device errors reported by the adapter driver.
  - LIP Resets: The number of loop initialization process resets reported by the adapter driver.
  - I/O Count - Total Number of I/Os reported by the adapter driver.
  - I/O Per second - Current number of I/Os per second.
  - Bytes Per Second - Current number of bytes processed by the adapter per second.
- Access to diagnostics
  - The read/write buffer test sends data to the target device using the SCSI Write Buffer command, and then reads the data back using the SCSI Read Buffer command, and compares for any errors.
  - The Link Status of the device is taken both before and after the read/write test, and the values compared. If errors occur, the test indicates a broken or unreliable link between the adapter and device.

- Alarms and indicators

  FAStT Check gives you the ability to set up polling for alarms for machines that are connected. The alarms will be kept in the FAStT Check Event Log. To access this, select **File->Event Log**. The entry in the Event Log describes the alarm, identifies the relevant machine and adapter, and provides a timestamp for the alarm.

3.3.1 Install Application

We now describe the steps necessary to install the FAStT Check utility.

1. Run the self-extracting file fastc_rl.exe.
2. During the install you will be asked if you wish to install Intel DMI 2.0 Service Provider SDK 1.10. as shown in Figure 3-9:

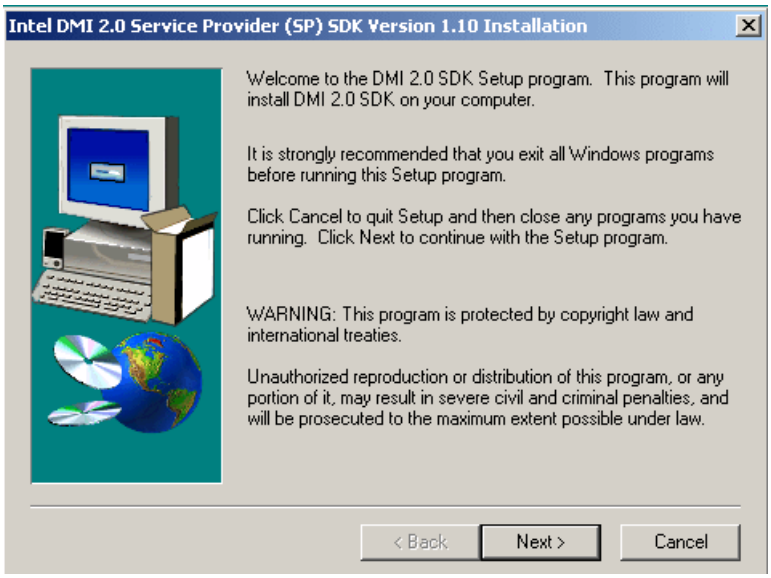

Figure 3-9 Installing FAStT Check

3. Click **Next** to be presented with check boxes for DMI 2.0 Service Provider and DMI 2.0 DCE Client. Check the boxes as shown in Figure 3-10:

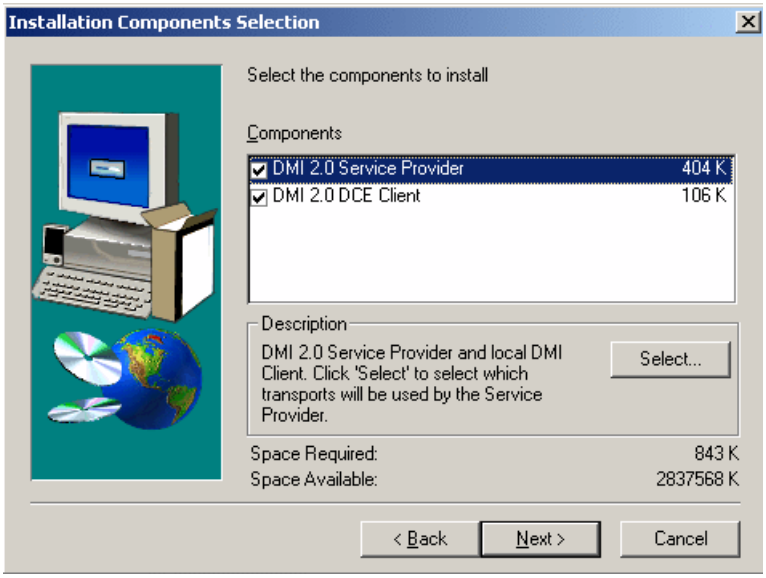

Figure 3-10 Install Components section

4. Click **Next** and then follow the displayed instructions to complete the installation.

3.3.2 Obtaining the World Wide Names of adapters

Since we will be using soft zoning to segment our SAN, we need to obtain the WWNs of the devices installed. This is because soft zoning converts WWNs to alias names.

1. Start the FAStT Check application to open the main window (shown in Figure 3-11):

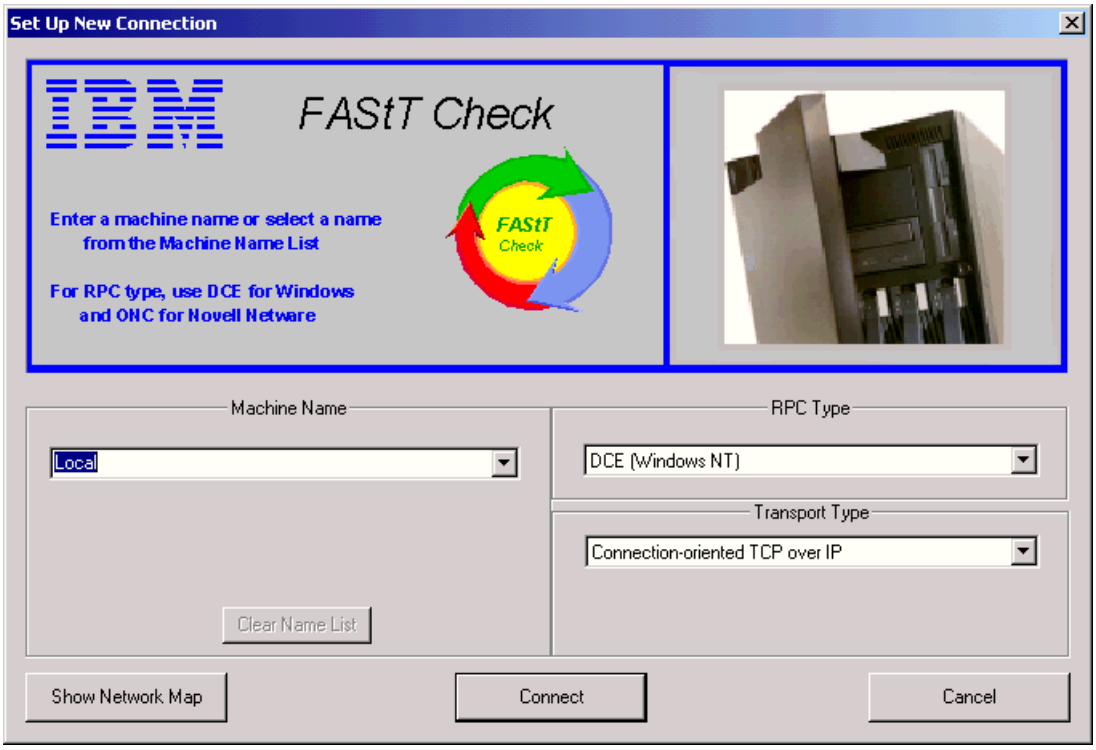

Figure 3-11 FAStT Check main window

2. Since this application has been installed locally, we do not need to enter the Machine Name. Click **Connect**.
3. The next window displays the FAStT Host Adapters that are installed in our xSeries server, as shown in Figure 3-12:

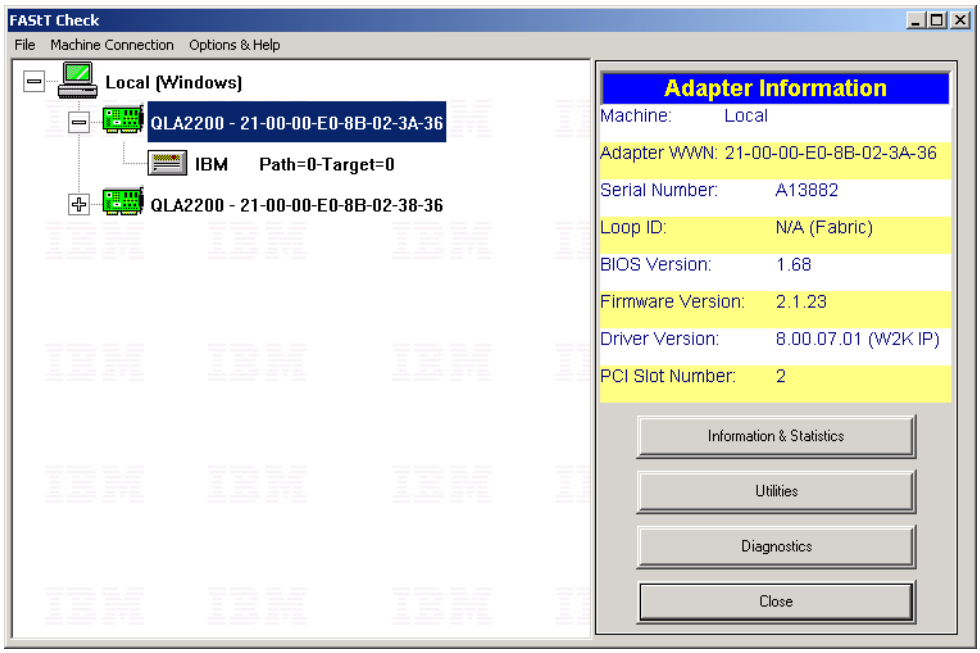

Figure 3-12 WWN and slot details

4. Record the WWN and slot details for each adapter for use when we configure software zoning.

4

Installation and configuration of IBM switches

This chapter describes the steps required to install and configure the IBM SAN switches for use. It covers the basics of zone configuration and managing the switch, and includes information about updating the IBM switch firmware.

4.1 Assigning an IP address to a Fibre Channel switch

Each switch requires its own unique IP address to enable communication with the device. Here we discuss how to assign addresses to your switches.

4.1.1 IBM 2109 Model S08

You assign an IP address to the 8-Port Fibre Channel Switch using a serial connection. If you prefer, you can attach the switch to your LAN using its Ethernet connection and configure the switch using a telnet session. However, to do so, you would have to change the network settings of the computer initiating the telnet session to the same network class and subnet as the switch. In this example, we use the serial port. For more information on configuring the IP address using the Ethernet port, please refer to your switch product documentation.

This switch has a factory default IP address, which is 10.77.77.77. To set the switch to a different IP address, you require a serial cable, a PC or laptop computer with an open serial port, and terminal emulation software. We recommend that you use HyperTerminal as your terminal emulator, since it is pre-installed on most Windows installations. Follow these steps to configure the IP address using a serial connection:

1. Attach the serial cable to the serial port of the switch and then to the PC or laptop computer's serial port, preferably COM1.
2. Start HyperTerminal from its shortcut or by running **`hypertrm.exe`**.
3. The application will start and present you with the Connection Description window. This is where you define a connection name, which is for convenience only and can be anything meaningful to you. We used `Switch`, as shown in Figure 4-1:

© Copyright IBM Corp. 2002

:

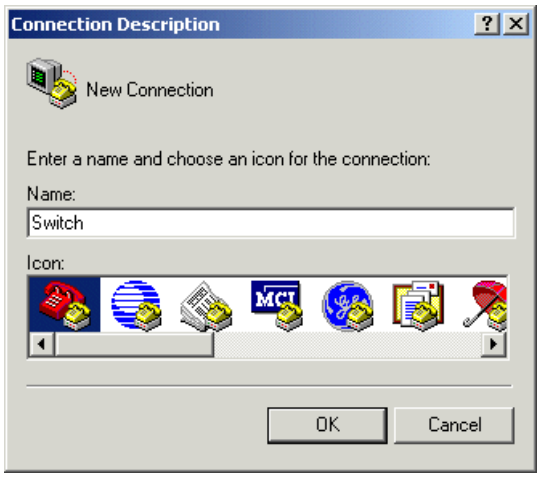

Figure 4-1 Connection Description window

4. Click **OK** to continue. The Connect To window is displayed. Here, you define how you wish to connect. We chose to connect using COM1, as shown in Figure 4-2:

Figure 4-2 Connect to window

5. If you are not using COM1, then you may specify a different port. Click **OK** to continue.
6. The COM1 Properties window is displayed, as shown in Figure 4-3:

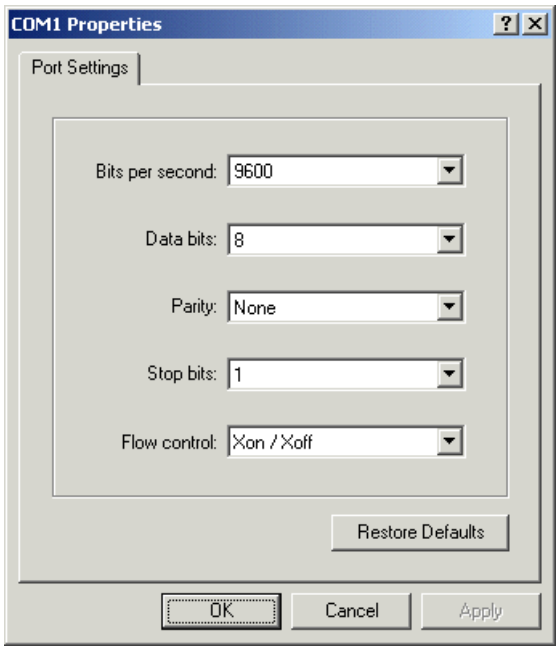

Figure 4-3 COM1 Properties window

7. In our lab, we chose 9600 bits per second and Xon/Xoff flow control. Click **OK** to continue
8. Press the Ctrl and Break keys at the same time.
9. Press the Enter key.
10. At the login prompt, type `admin`. This is the default username.
11. At the password prompt type `password`. This is the default password for admin.
12. At the prompt run the command **`ipAddrSet`**. This command steps you through the settings for the switch, one at a time. If you do not wish to change a particular setting, press Enter. Otherwise, enter the new value and then press Enter.

    Example 4-1 shows the output we saw in our lab. We changed the IP address and the related subnet mask:

Example 4-1 ipAddrSet command

```
IBM_2109_FC_Switch:admin> ipAddrSet
Ethernet IP Address [9.9.9.67]: 172.16.4.50
Ethernet Subnetmask [255.255.255.0]: 255.255.252.0
Fibre Channel IP Address [none]:
Fibre Channel Subnetmask [none]:
Gateway Address [192.168.90.1]: 172.16.4.1
Set IP addresses now?
[y = set now, n = next reboot]: y
```

13. As you can see in the example, once you have stepped though all the values, the switch prompts you to decide whether you wish to set the IP address now or at the next reboot. Press y and then the Enter key.
14. At the prompt run the command **`ipAddrShow`** to output the current configuration as shown in Example 4-2:

Example 4-2 ipAddrShow command

```
IBM_2109_FC_Switch:admin> ipAddrShow
Ethernet IP Address: 172.16.4.50
Ethernet Subnetmask: 255.255.252.0
Fibre Channel IP Address: none
Fibre Channel Subnetmask: none
Gateway Address: 172.16.4.1
IBM_2109_FC_Switch:admin>
```

15. At the prompt, enter the command **Reboot**. This will reboot the switch.

You have now configured the IP address of the switch. It is now ready to be attached using the Ethernet port to the LAN.

4.1.2 IBM 2109 Model S16

Just as for the 8-port switch, this model has a factory default IP address of 10.77.77.77, and, if you prefer, you can attach the switch to your LAN using the Ethernet connection and configure the switch using a telnet session. To do so, you would have to change the network settings of the computer initiating the telnet session to the same network class and subnet as the switch. In this example, however, we set up the IP address using the keys and LCD display available on the 16-port switch. For more information on configuring the IP address using the Ethernet port, please refer to your switch product documentation.

The following steps should be followed when using the LCD to configure network settings:

1. Make sure the switch is powered up. Locate the buttons located on the front of the switch.
2. Click the ▲ button until you reach the Configuration menu.
3. Click the ▶ button, this takes you into the Configuration menu options.
4. Click the ▲ button until you reach Ethernet IP address.
5. Click the ▶ button, this takes you into the Ethernet IP address option.
6. Click the ◀ button to move between locations within the IP address. Click the ▲ button to change the values.
7. Click the ▶ button to commit the changes. The switch will respond Yes or No. Click ◀ for Yes or ▶ for No.
8. Click the ▲ button until you reach Ethernet Subnet Mask.
9. Click the ▶ button, this takes you into the Ethernet Subnet Mask option.
10. Click the ◀ button to move between locations within the Subnet Mask. Click the ▲ button to change the values.
11. Click the ▶ button to commit the changes. The switch will respond Yes or No. Click ◀ for Yes or ▶ for No.
12. Click the ▲ button until you reach Gateway Address.
13. Click the ▶ button, this will take you into the Gateway address option.
14. Click the ◀ button to move between locations within the IP address. Click the ▲ button to change the values.
15. Click the ▶ button to commit the changes. The switch will respond Yes or No. Click ◀ for Yes or ▶ for No.

You have now configured the IP address of the switch, and it is ready to be attached to the LAN using the Ethernet port.

4.2 Install Java

The preferred management tool for the SAN switches is accessed using a Web browser. The Web browser must have Java enabled for proper operation.

The Java 2 Version 1.3.1_01 code is available for download from the following site:

http://java.sun.com/j2se/1.3/download-windows.html

After downloading, execute the program and follow the instructions to complete installation.

4.3 Setting a unique switch name

To make sure that the switches are unique on the SAN, we have to give each one its own switch name.

1. Open your Web browser and enter the IP address of the switch, which we set in 4.1, “Assigning an IP address to a Fibre Channel switch” on page 29. You will then see the main switch view. The window is shown in Figure 4-4:

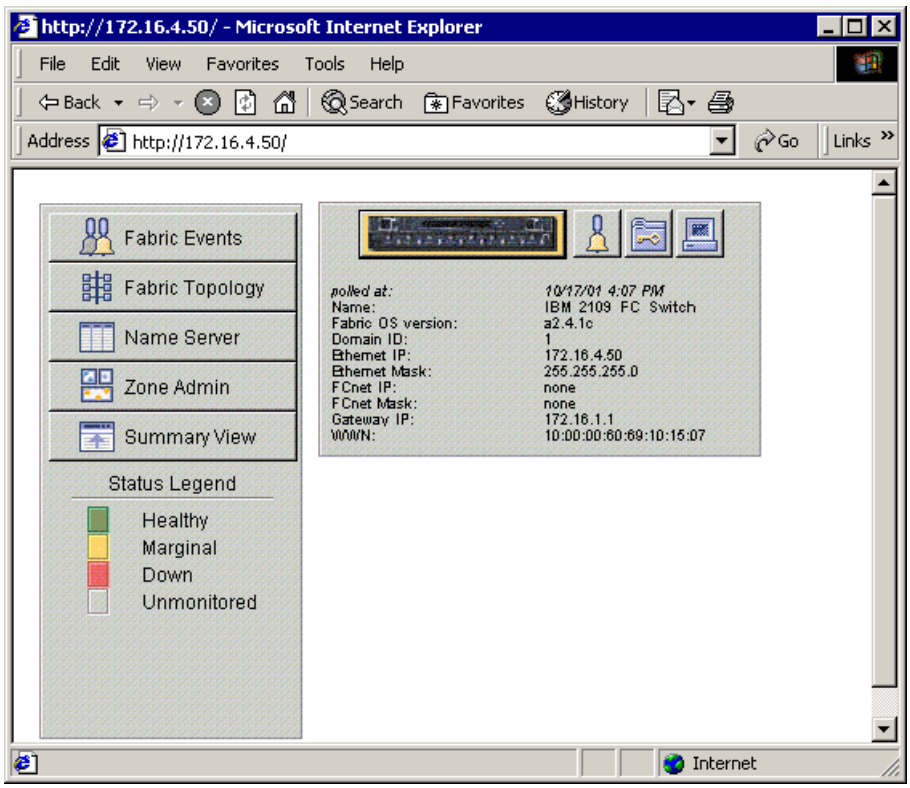

Figure 4-4 Switch main window

2. Click the **Administration** button  to be presented with a login window as shown in Figure 4-5:

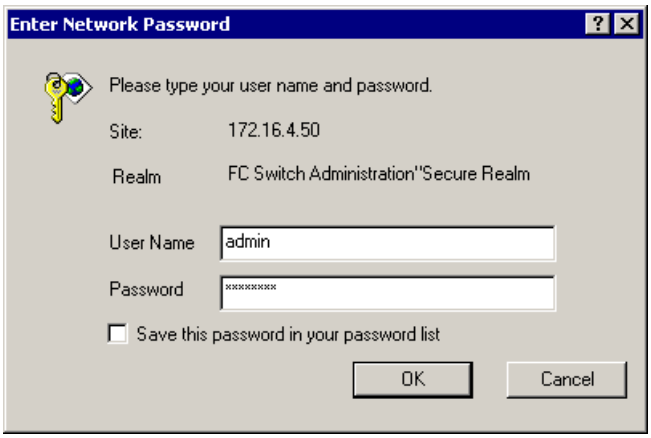

Figure 4-5 Switch Login.

3. Enter the default user ID and password, which are `admin` and `password` respectively, and click **OK** to display the switch administration window, shown in Figure 4-6:

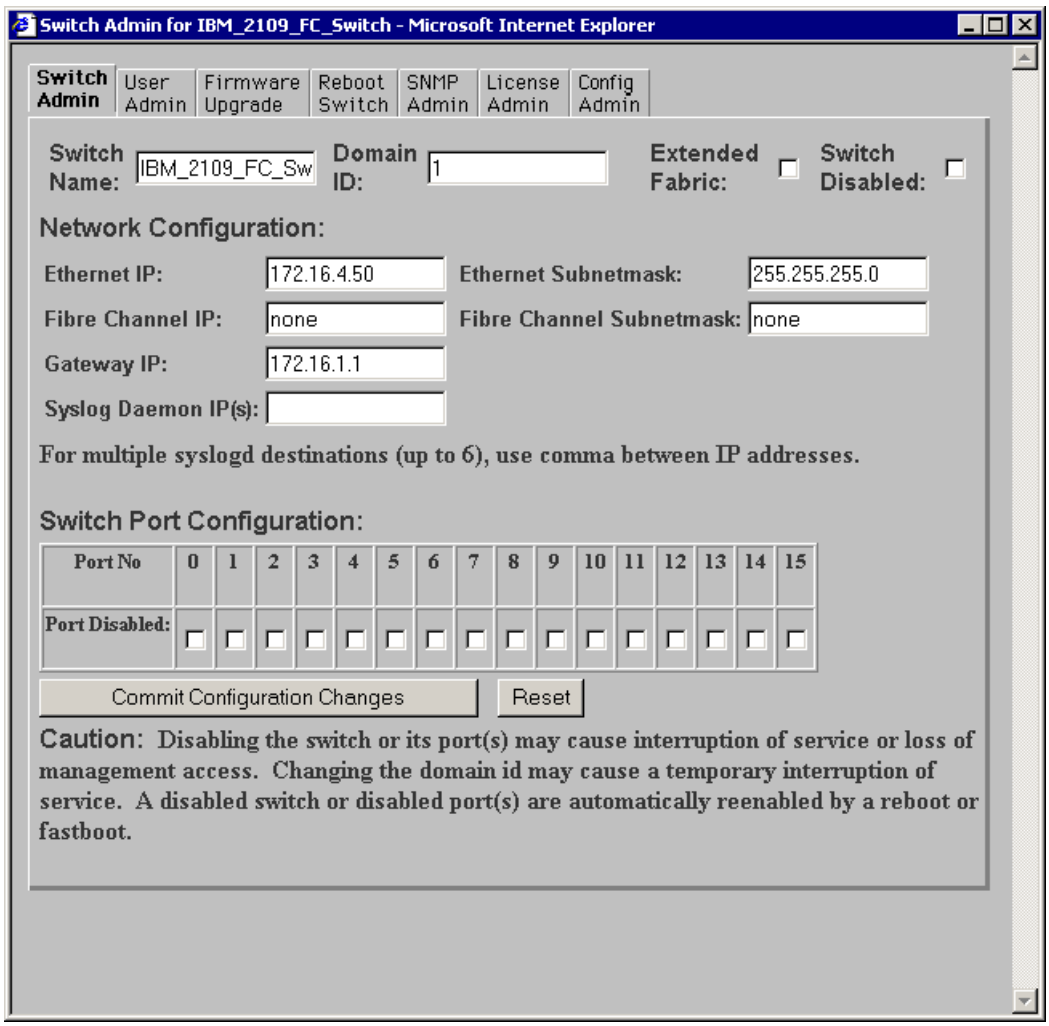

Figure 4-6 Switch administration window

4. Input a new name in the Switch Name field (we used `ITSO_Switch_1`), then click **Commit Configuration Changes**. You can also select the User Admin tab at this point and change the password if you wish.
5. The switch now has a new name. We did not change the Domain Id for this switch because it is the only switch on the SAN.

4.4 Recording World Wide Names

We strongly recommend that you document all connections in your SAN. This will allow you to configure and administer your SAN more efficiently. One of the best sources of information is your IBM switch. We used this to obtain the WWN's so that we could complete our zoning.

From the main window, shown in Figure 4-4, click the **Name Server** button. This will list the Fibre Channel devices connected to the switch, as shown in Figure 4-7:

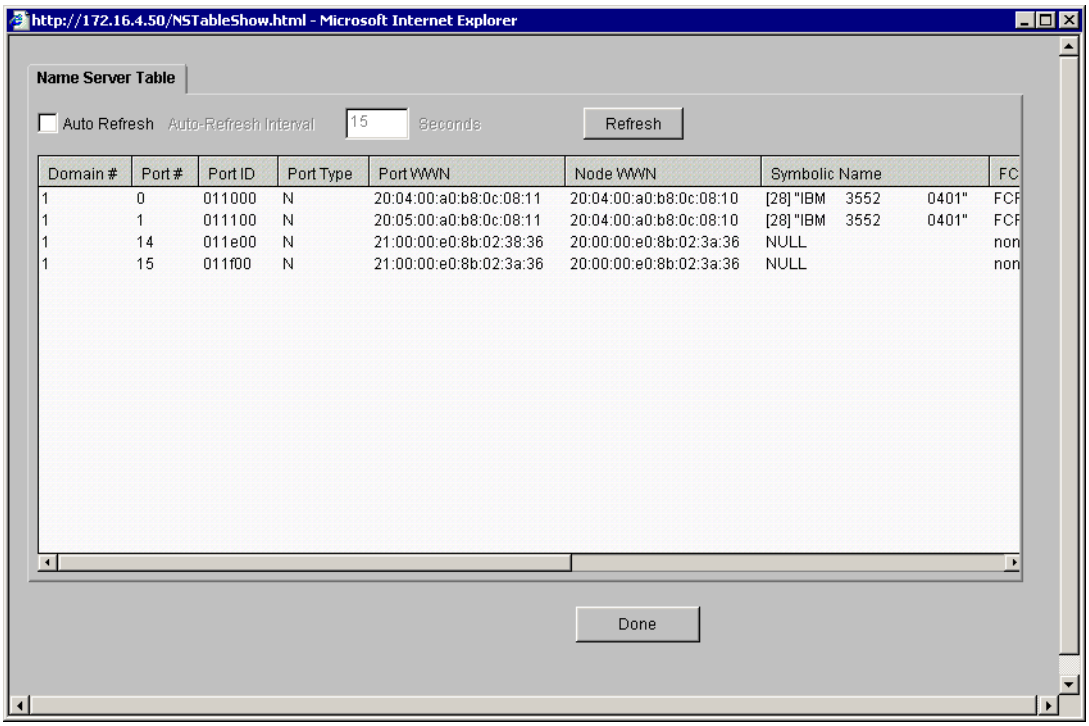

Figure 4-7 Name Server Table

4.5 Switch firmware upgrade

To ensure that all the code levels throughout the SAN are at the latest levels, it is important to update the switch firmware.

1. Create a directory on your local machine into which you can download your firmware.
2. Copy the files rshd.exe, cat.exe (from the rsh.zip file) and a2.4.1c (or the latest firmware) to the new directory.
3. Execute the file rshd.exe (rsh daemon). This remote shell daemon needs to run in the background during a firmware update, as shown in Figure 4-8:

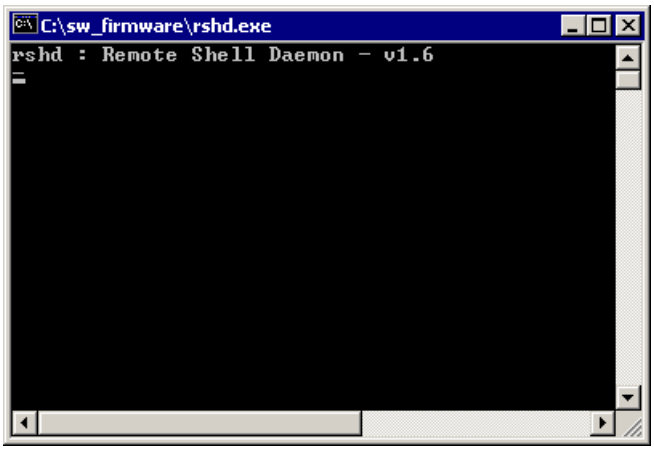

Figure 4-8 Remote Shell daemon

4. Open your Web browser and connect to the switch as before. Click the **Administration** button, log in, and select the **Firmware Upgrade** tab as shown in Figure 4-9:

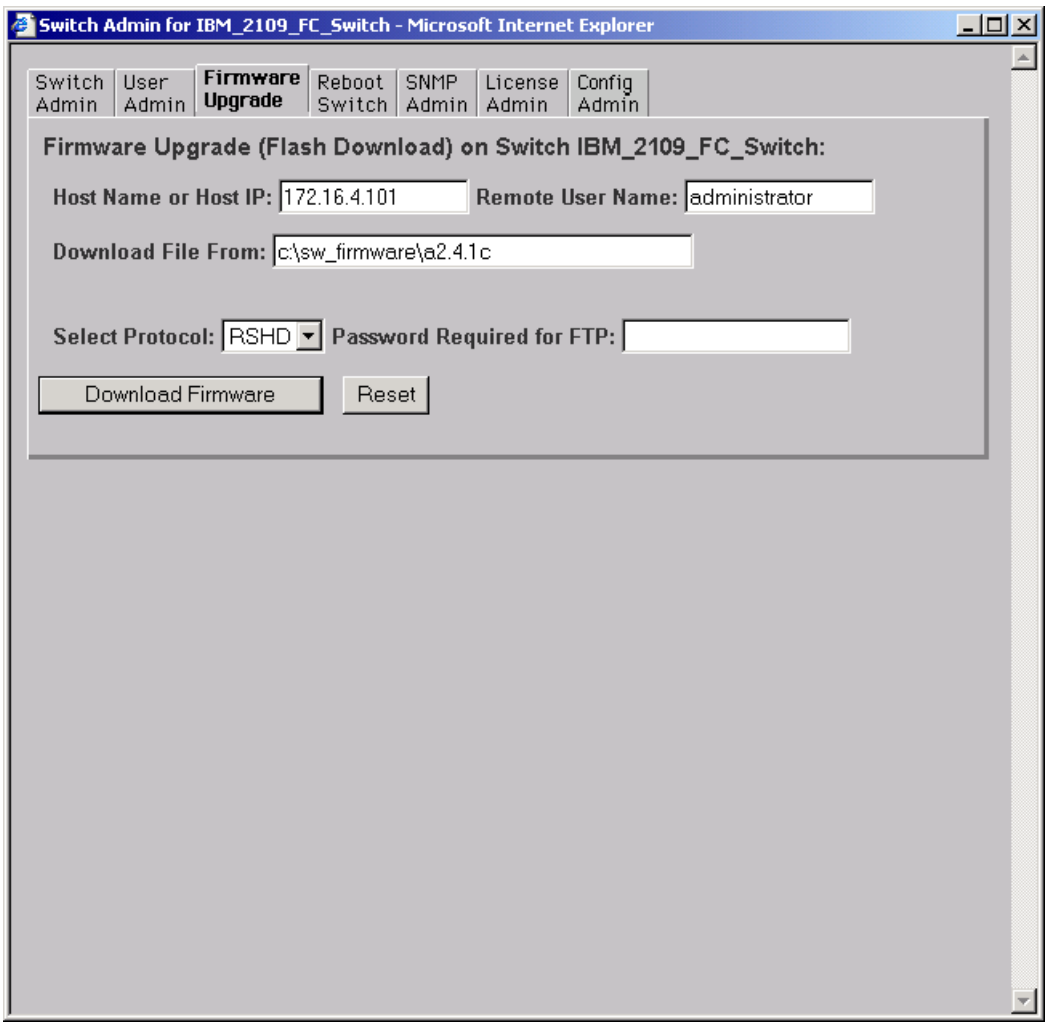

Figure 4-9 Updating the switch firmware

5. Enter the host name or IP address of the local machine and the Remote User Name (this is the user who is running the remote shell). In our example, this was 172.16.4.101 and Administrator.
6. In the Download File From field, enter the path and file name of the firmware upgrade file. Set the Select Protocol field to RSHD. Click **Download Firmware**.
7. Once the download has completed you need to reboot the switch for the new firmware to take effect. When the pop-up appears, click **Proceed to reboot the switch**.
8. The Reboot Switch tab is displayed as shown in Figure 4-10:

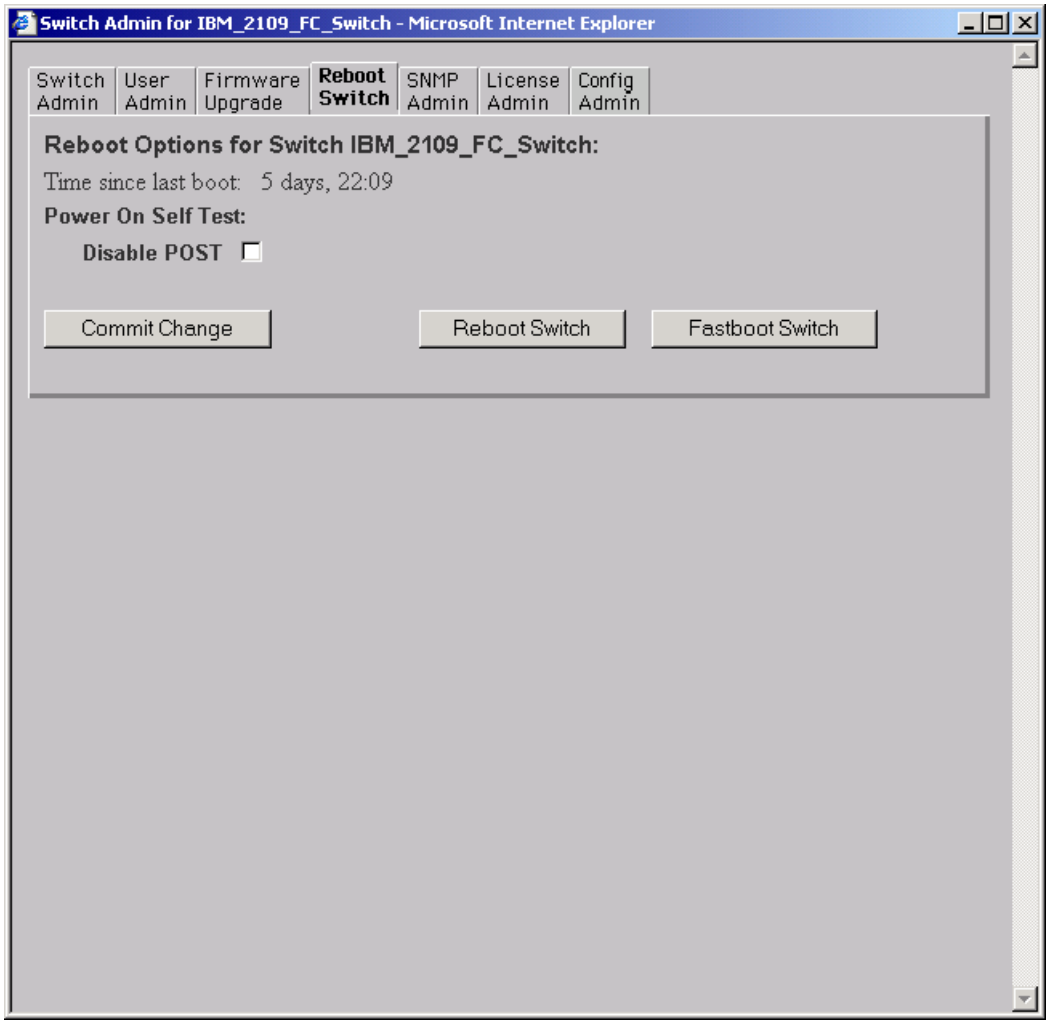

Figure 4-10 Reboot switch

9. Click **Fastboot Switch** to omit the diagnostic check and thus allow the switch to boot more quickly.
10. The new firmware is now in effect. You can confirm this by checking the Fabric OS version on the main switch window (see Figure 4-4 on page 33).

4.6 Zone configuration

1. The final step in configuring the SAN fabric is to zone the switches. This allows the SAN to be segmented. The reasons you may decide to segment your SAN may be for functional,

security or heterogeneous operating system issues. We opted for software zoning, as described in 2.2.3, “Zoning” on page 14. Briefly, the reasons for choosing software zoning were twofold. First, we could use a simple naming convention, and, second, the cabling can be moved across to other switches for redundancy without disturbing the zone configuration.

4.6.1 Creating aliases

To simplify administration, we first give each device on the SAN a more meaningful name. Here are the steps:

1. At the switch main window (Figure 4-4 on page 33), select **Zone Admin**.
2. We previously recorded all the WWNs (4.4, “Recording World Wide Names” on page 35) and we can now make them more memorable by creating aliases.
3. Click **Create Alias** to be asked to enter a name, as shown in Figure 4-11

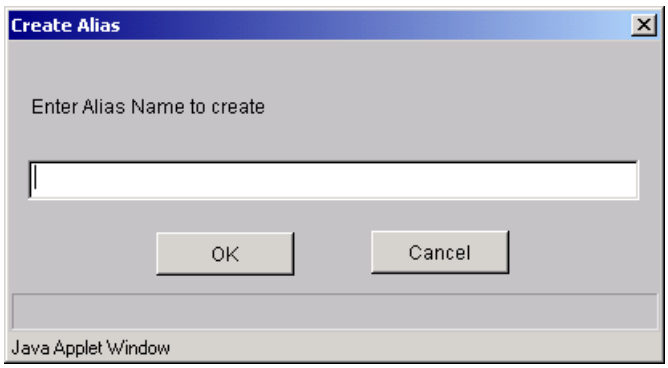

Figure 4-11 Create Alias window

4. Enter the Alias name and click **OK**.
5. Highlight the device’s WWN in the left pane and click **Add member**, then click **Apply**.

> **Note:** Do not choose **Add Host**, but always choose **Add member** when creating aliases.

Figure 4-12 shows an alias being created for the FAstT Host Bus Adapters:

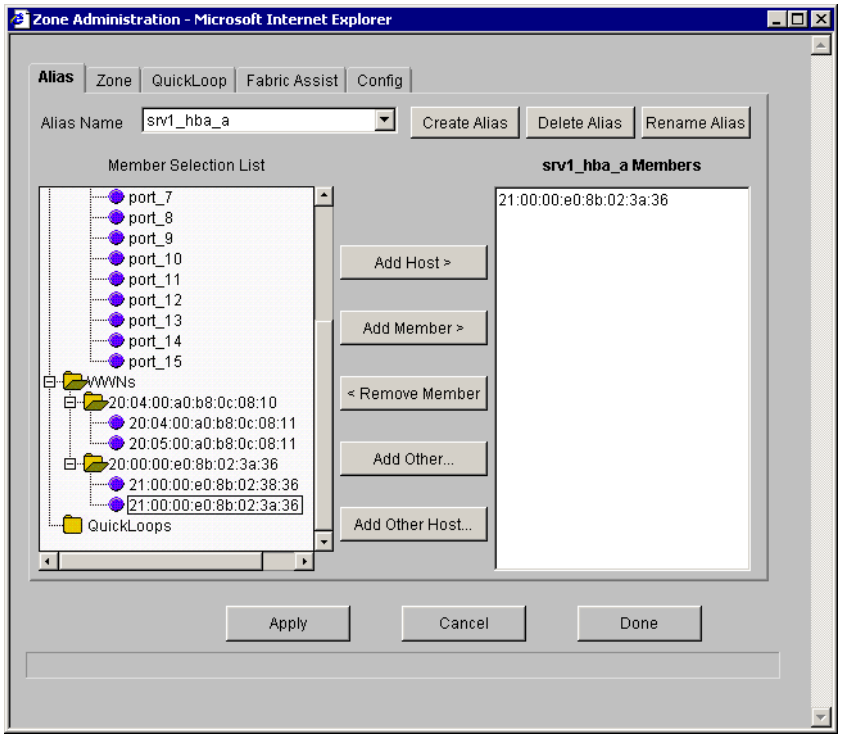

Figure 4-12 Create alias for FAStT Host Bus Adapter

Figure 4-13 shows an alias being created for the FAStT500 Controller A:

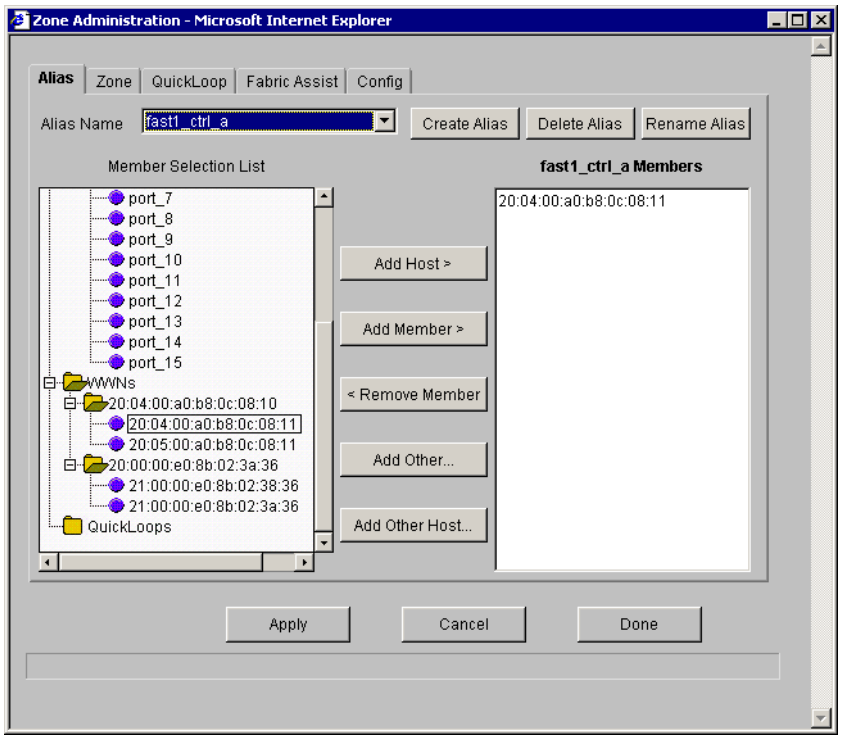

Figure 4-13 Create alias for FAStT500 RAID Controller A

4.6.2 Creating zones

As we have multiple FAStT Host Bus Adapters installed, this means that we have multiple paths to the disk drives. In Windows this results in the drives appearing twice in Disk Administrator. To correct this, we zone the switch so that each card has only one path to a FAStT500 RAID controller. Here are the steps:

1. In the Zone Administration window, select the **Zone** tab.
2. Click **Create Zone**. A pop-up window appears in which you enter the zone name (see Figure 4-14):

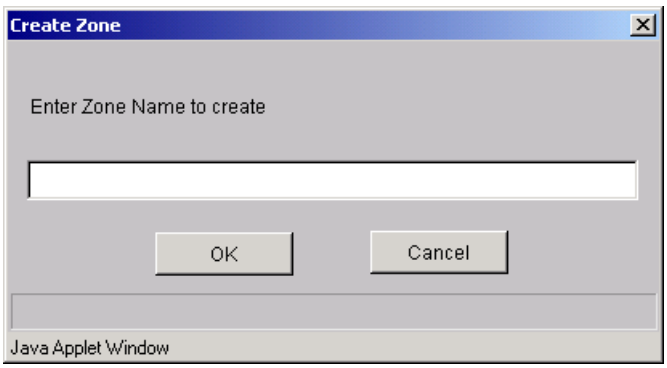

Figure 4-14 Create zone

3. Enter the name of the zone. We entered `fast1_zone_a`. Click **OK**.
4. Highlight the alias name in the left pane and click **Add Member** to add the alias to the zone.
5. Once all members of the zone have been added, click **Apply**.

In Figure 4-15 you can see that we have created a zone called fast1_zone_a, which includes the host bus adapter of server1 and controller A of the FAStT500:

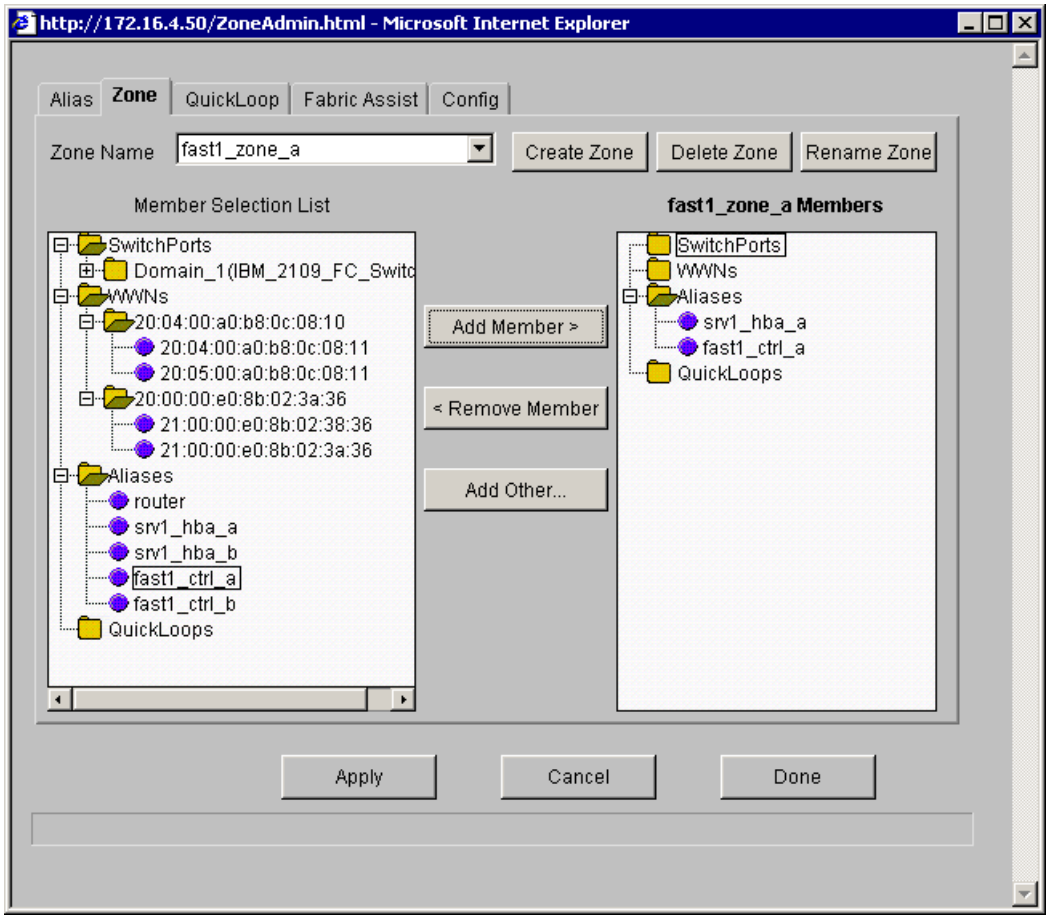

Figure 4-15 Zone configuration

6. Repeat steps 2 to 5 for the second FAStT HBA and controller B of the FAStT500 RAID Controller.

4.6.3 Creating and enabling a zone configuration

Now that you have created the zones, you have to create and enable the configuration. It is possible to have multiple configurations defined, but only one can be enabled at any time. This flexibility allows SAN administrators to create test zones and to split their SAN environments to respond to changing requirements. Here is how a configuration is enabled:

1. Select the **Config** tab in the administration window.
2. Click **Create Cfg** to display the pop-up shown in Figure 4-16:

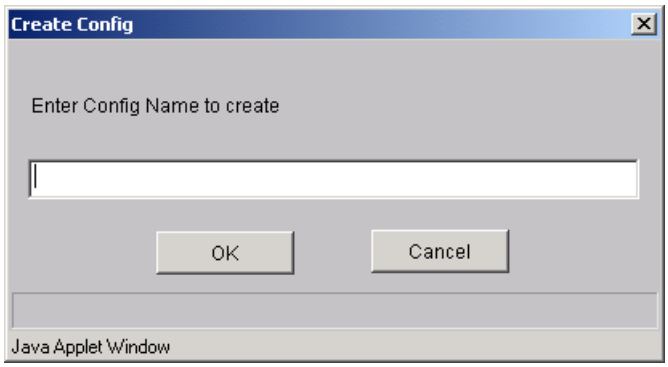

Figure 4-16 Create Config window

3. Enter a configuration name. We called our configuration `ITSO`.
4. Highlight a zone in the left pane and click **Add member** to add it to the configuration. Repeat for all member zones. An example of our SAN configuration is shown in Figure 4-17:

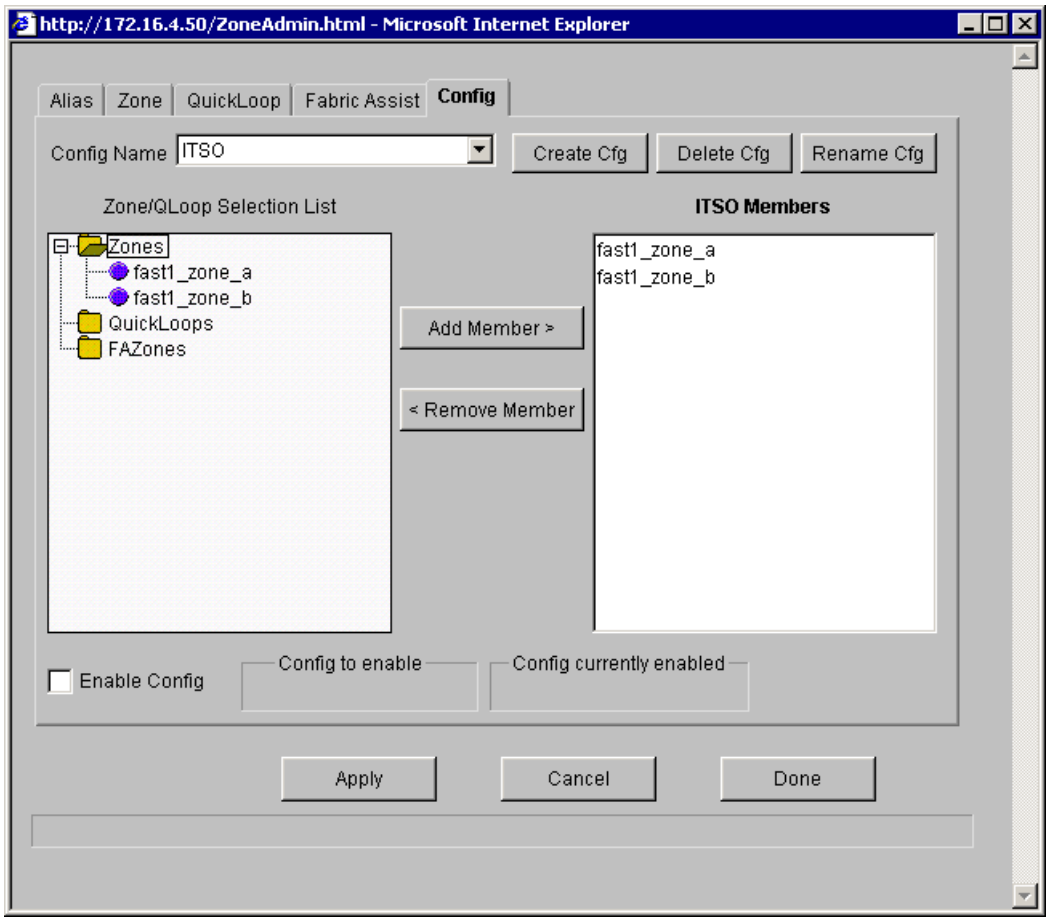

Figure 4-17 Zone configuration

5. Now that you have a configuration, it needs to be enabled. Check **Enable Config** as shown in Figure 4-18:

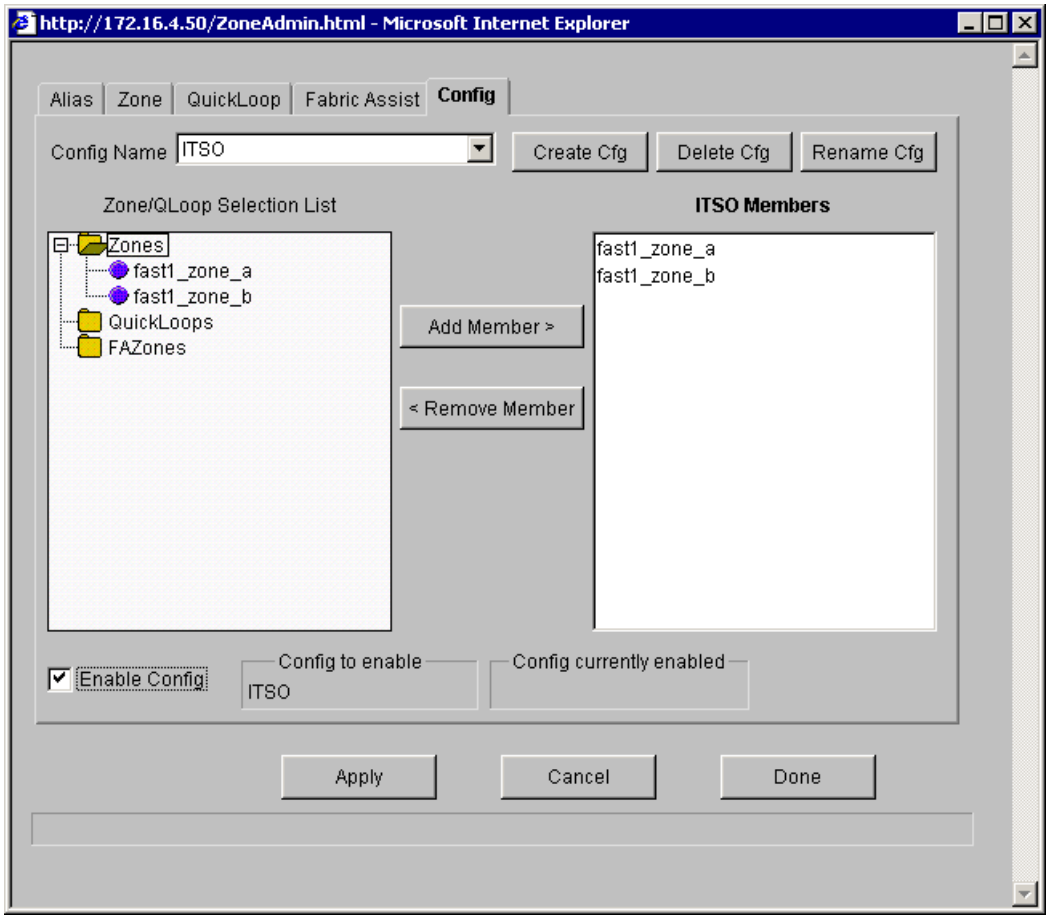

Figure 4-18 SAN configuration enabled

6. Click **Apply** and then **Done**. You have now completed the switch setup.

To confirm that your zones are working correctly, go to the main menu of the switch and select the **Name Server Table** option. This will give you details of the devices in the SAN. If you scroll all the way to the right you can see the zones. If the zone has an asterisk then the zone is enabled, as shown in Figure 4-19:

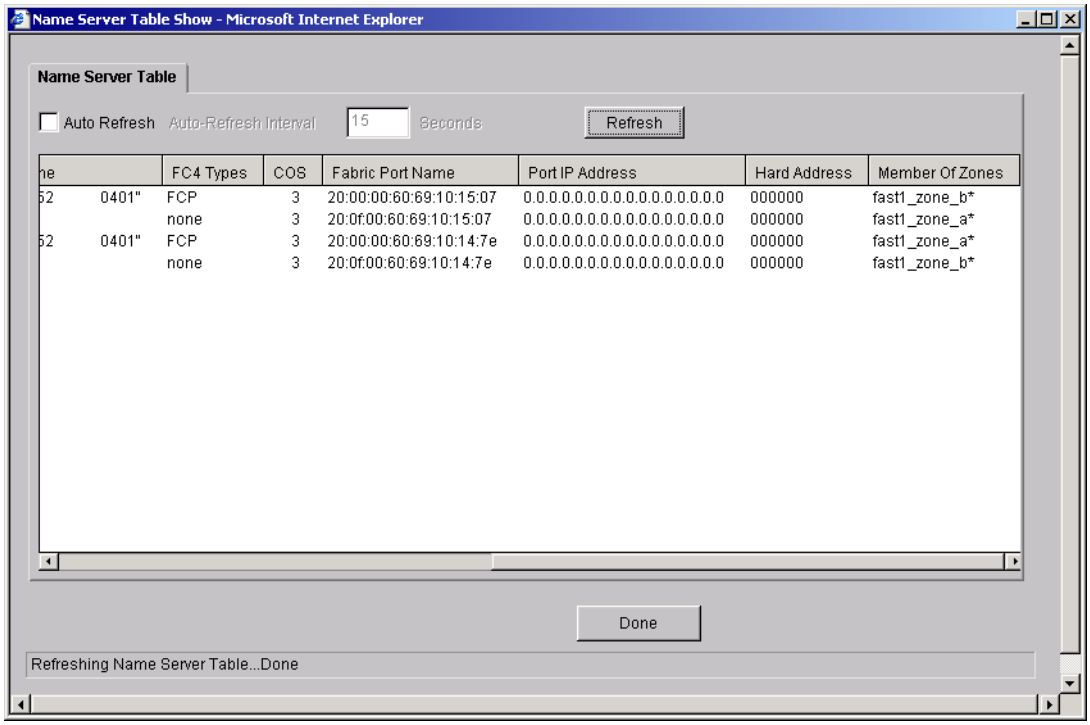

Figure 4-19 Confirmation of zone configuration

5

FAStT500 and FAStT Storage Manager Setup

In this chapter we examine the FAStT500 RAID Controller and the FAStT Storage Manager V7.10 software. We configure the storage server and software using a step-by-step approach and suggest the best practices that should be followed.

As it is not intended that this Redpaper serve as a manual, a cookbook approach is taken, keeping detailed explanations to a minimum. If you feel that you require more information on a specific topic, please refer back to the product documentation. Our intent is to provide you with the information required to get your SAN up and running correctly as soon as possible, avoiding the need to sift through hundreds of pages of documentation.

5.1 FAStT500 setup

In this section we discuss how to set up the FAStT500 RAID Controller so that it can be configured and managed using the Storage Manager software.

5.1.1 Assigning an IP address to a FAStT500 controller

Here we discuss how to assign an IP address to the FAStT500 controllers, using a serial connection to perform direct-attachment management.

Important: The serial connection should be used only to perform these commands, unless otherwise instructed to do so by IBM Support.

To set the IP address of either controller, you will need a null modem cable, PC or laptop computer with an open serial port, and terminal emulation software. We recommend that you use HyperTerminal as your terminal emulator as it is pre-installed on most Windows systems.

1. Attach the null modem cable to the serial port of the FAStT500 Controller A and then to the PC or laptop computer's serial port, preferably COM1.
2. Start HyperTerminal from its shortcut or by running **`hypertrm.exe`**.

© Copyright IBM Corp. 2002

3. The application starts, and the Connection Description window is displayed. This is where you define a connection name, such as `FAStT500`, which we used, as shown in Figure 5-1:

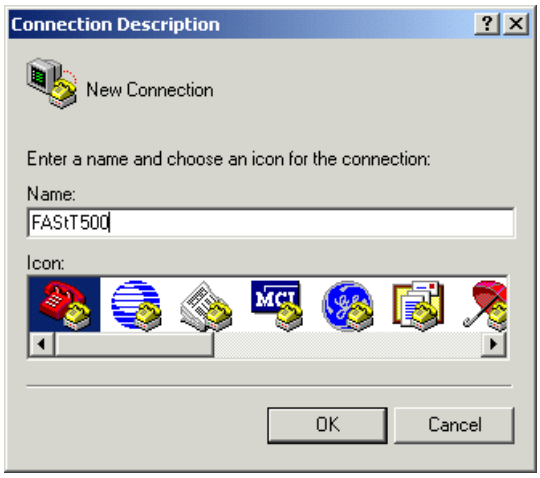

Figure 5-1 Connection Description window

4. Click **OK** to display the Connect To window as shown in Figure 5-2:

Figure 5-2 Connect To window

5. This window lets you define how you wish to connect. If you are not using COM1, specify the COM port to which you connected the null modem cable. Click **OK** to display the COM port Properties window as shown in Figure 5-3:

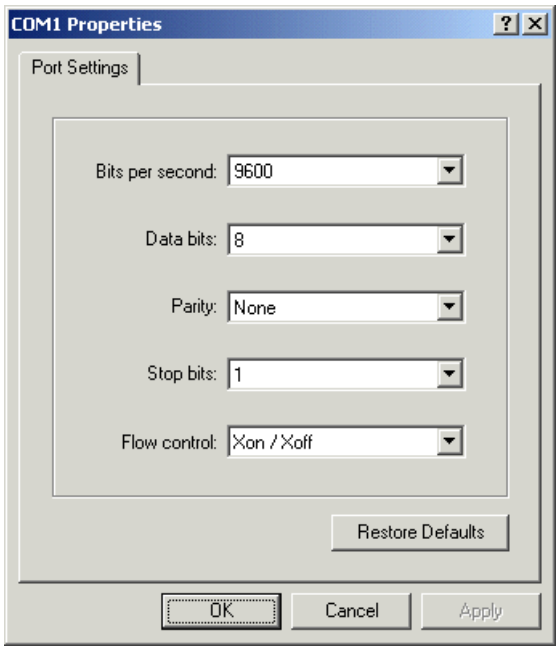

Figure 5-3 COM1 Properties window

6. Set the link speed (bits per second) to 9600 and flow control to Xon/Xoff. Leave data bits, parity, and stop bits at their default settings of 8, None, and 1 respectively. Click **OK** to continue.
7. Press the Ctrl and Break keys at the same time.
8. Press the Esc key to display the FAStT500 Shell login window as shown in Figure 5-4:

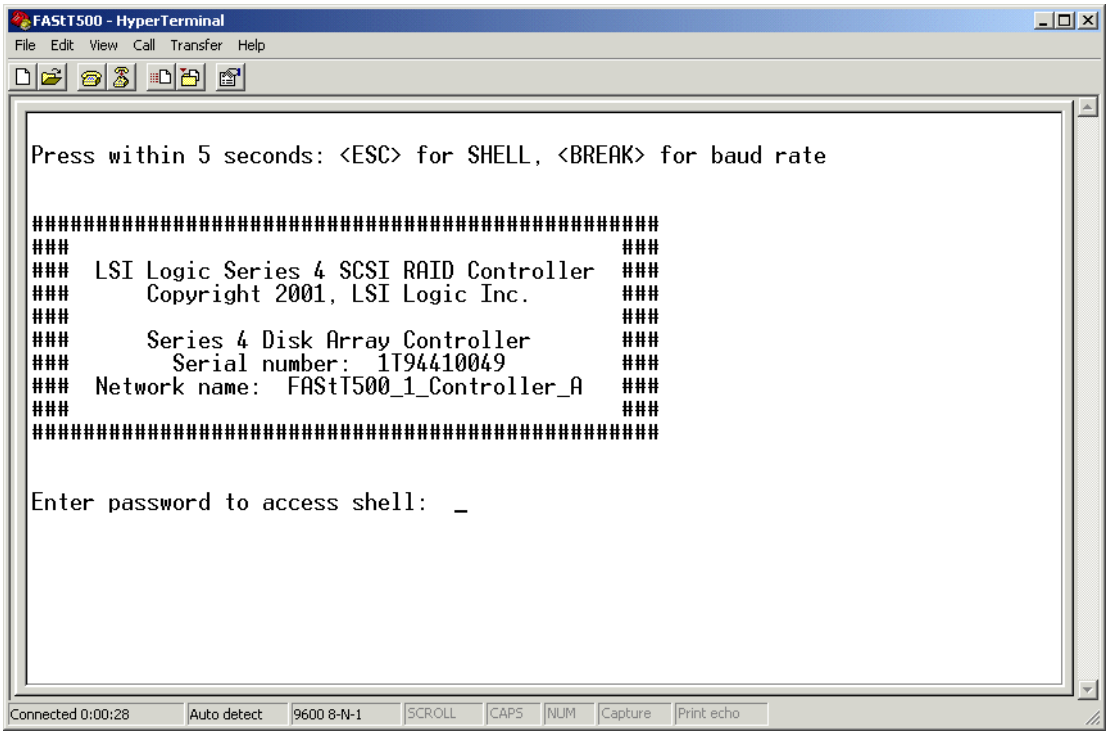

Figure 5-4 FAStT500 password prompt

9. At the password prompt enter `infiniti`, which is the default password for a new FAStT500 controller.

10. At the prompt run the command **`netCfgSet`**.

    Press the Enter key, accepting the current value for each setting, until you reach the item that you wish to change. Type in the new value and press Enter to move to the next setting.

    The settings we changed are indicated in bold type in Example 5-1:

Example 5-1 netCfgSet output

```
-> netCfgSet


    '.' = clear field;  '-' = to previous field;
    '+' = next interface;  ^D = quit (keep changes)

==== NETWORK CONFIGURATION: ALL INTERFACES ====
Network Init Flags    :  0x01
Network Mgmt Timeout  :  30
Network Route #1      :  dest=0.0.0.0
RAIDMGR Server #1     :  0.0.0.0
Network Manager #1    :  0.0.0.0
Startup Script        :
Shell Password        :

==== NETWORK CONFIGURATION: Ethernet ====
My MAC Address        :  00:a0:b8:06:64:89
My Host Name          :  target FAStT500_1_Controller_A
My IP Address         :  10.10.10.1 172.16.4.70
Server Host Name      :  host
Server IP Address     :  0.0.0.0
Gateway IP Address    :  0.0.0.0 172.16.4.1
Subnet Mask           :  255.255.255.0 255.255.252.0
User Name             :  guest
User Password         :
NFS Root Path         :
NFS Group ID Number   :  0
NFS User ID Number    :  0
value = 0 = 0x0
->
```

11. At the prompt run the command **`netCfgShow`**, which allows you to check that your configuration is correct (see Example 5-2):

Example 5-2 netCfgShow output

```
-> netCfgShow

==== NETWORK CONFIGURATION: ALL INTERFACES ====
Network Init Flags    :  0x01
Network Mgmt Timeout  :  30
Startup Script        :
Shell Password        :
==== NETWORK CONFIGURATION: Ethernet ====
Interface Name        :  dse0
My MAC Address        :  00:a0:b8:06:64:89
My Host Name          :  FAStT500_1_Controller_A
My IP Address         :  172.16.4.70
Server Host Name      :  host
Server IP Address     :  0.0.0.0
```

```
Gateway IP Address    :  172.16.4.1
Subnet Mask           :  255.255.252.0
User Name             :  guest
User Password         :
NFS Root Path         :
NFS Group ID Number   :  0
NFS User ID Number    :  0
value = 0 = 0x0
->
```

12. At the prompt run the command **sysReboot**, which reboots the connected controller.

Perform the same configuration procedure for the FAStT500 Controller B, bearing in mind that its host name and IP address must not be the same as for controller A. The subnet mask and gateway IP address should be the same for both controllers.

After performing this procedure on both controllers, they are ready to be connected to your local area network.

Tip: We recommend that you use an isolated or firewall-protected segment of your network that is only accessible by a management workstation and not by the rest of your LAN, as is shown in Figure 5-5:

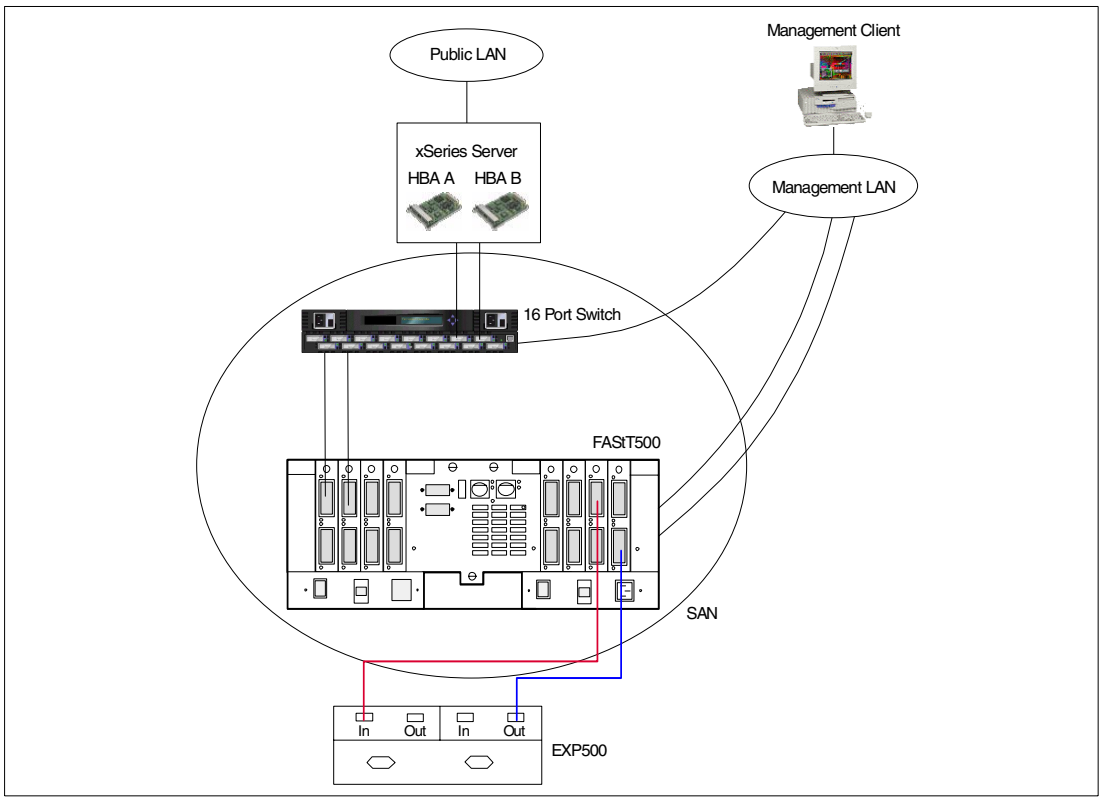

Figure 5-5 LAN segment for SAN management

13. Using the **ping** command, ping both controller IP addresses to check that they are connected to the LAN before moving on to installing FAStT Storage Manager.

5.2 FAStT Storage Manager 7.10

In this section we discuss the installation of the FAStT Storage Manager software. Specifically, we install the following items:

- The Redundant Dual Active Controller (RDAC) driver that provides redundant paths to your storage server
- The server side utilities
- The client used to manage storage servers

We do not cover use of the Host Agent, since we feel that it is preferable to manage a storage subsystem using the direct-attach (out-of-band) option using Ethernet rather than using the host-attached (in-band) option over fiber.

Host Agent allows you to manage the storage subsystem using the fiber connection. This means you are limited to installing the software on machines physically attached to the SAN, although you can manage the storage subsystem remotely by talking to a machine running Host Agent.

We believe that you will almost always want to manage the Storage Server remotely, and it is simpler if you communicate directly with the Storage Server. You might prefer to use the Host Agent if you do not want to enable IP on your Storage Server for security reasons, but we believe it is preferable to connect your SAN devices to a separate LAN segment, that is either physically isolated or firewalled from other networks.

Using Host Agent offers no added functionality and adds complexity. If, for example, your Host Agent server is unavailable, you cannot manage your Storage Server. To avoid this, you have to install Host Agent on multiple machines.

If you still feel that you would like to use Host Agent, please refer to the Storage Manager 7.10 product documentation.

When installing the Storage Manager software and components, please remember to be logged in under the operating system as administrator or using an account with administrative rights for the machine on which you wish to install the code.

5.2.1 Installing the Storage Manager RDAC

The RDAC, or multi-path, driver should be installed on each server attached to your FAStT Storage Server SAN, and that has dual Host Bus Adapters installed. The RDAC does not need to be installed on servers with only one Host Bus Adapter. We do not recommend using only one Host Bus Adapter, however, as this creates a single point of failure between the host and the storage server. At the time of writing this paper, the RDAC was only available for Windows.

Tip: Always use two Host Bus Adapters in a server and install the RDAC.

Follow these steps to install the RDAC driver:

1. Run the downloaded installation package. In our example, the file name was SM7rdac-2000-07109503.exe. The Installation Folder window is displayed, as shown in Figure 5-6:

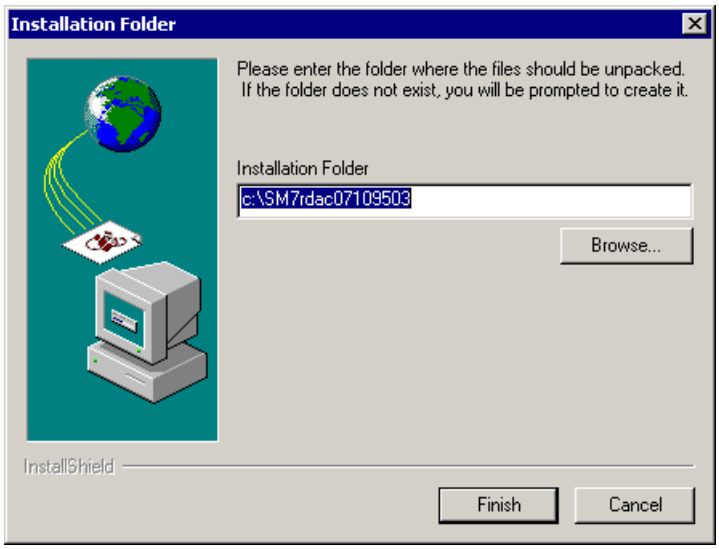

Figure 5-6 Installation Folder window

2. Enter the path to the folder where you wish to decompress the installation package and click **Finish**.
3. If the specified folder does not exist a window will follow. Click **Yes** and the folder will be created for you.
4. Once the package has been decompressed, a message will inform you that the package has been delivered sucessfully. Click **OK** to continue.
5. Using My Computer or Windows Explorer, locate the folder in which the decompressed RDAC package was stored. To begin the installation of the RDAC, run **`setup.exe`**.
6. An information message is displayed. Click **OK** to continue.
7. The IBM FAStT Storage Manager 7 Welcome window is displayed. Click **Next** to continue.
8. A status window is displayed, indicating the progress of the installation.
9. Once the installation is complete, an InstallShield Wizard Complete window appears. Select **Yes, I want to restart my computer now**, and click **Finish** to end the installation and reboot the server.
10. Once the server has restarted and you have logged on, the Found New Hardware Wizard is started, indicating that the RDAC Virtual Adapter has been found (see Figure 5-7):

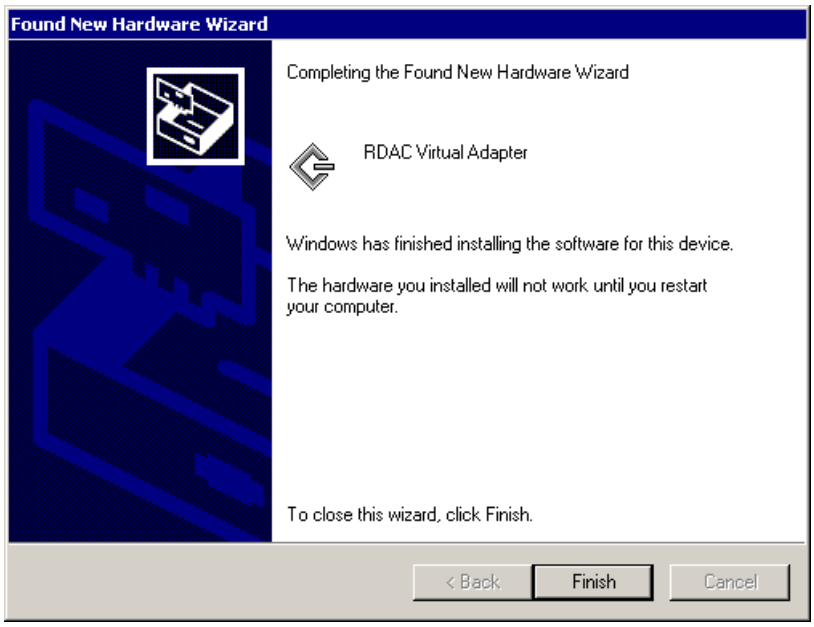

Figure 5-7 Found New Hardware Wizard window

11. Click **Finish**, and the System Settings Change message will follow. Click **Yes** to restart the server.

The RDAC has now been successfully installed.

5.2.2 Installing the Storage Manager utilities

We recommend that you install these utilities on all servers attached to your FAStT Storage Server SAN. They are used to perform tasks such as the hot-add of disks and are useful when troubleshooting your SAN. This section covers the installation of these utilities. In 5.3.13, “Using the Storage Manager utilities” on page 69, we show you how to use the utilities to hot-add a disk and to troubleshoot.

Follow these steps to install the utilities:

1. Run the downloaded installation package. In our example the file name was SM7util-WIN32-07102501.exe. The Installation Folder window is displayed.
2. Enter the path to the folder where you wish to decompress the installation package and click **Finish**.
3. If the specified folder does not exist, a window will follow. Click **Yes** and the folder will be created for you.
4. Once the package has been decompressed, a message will inform you that the package has been delivered sucessfully. Click **OK** to continue.
5. Using My Computer or Windows Explorer, locate the folder in which the decompressed package was stored. To begin the installation of the utilities, run **`setup.exe`**.
6. The IBM FAStT Storage Manager 7 Welcome window is displayed. Click **Next** to continue.
7. The Choose Destination Location window appears. Specify where the utilities are to be installed or accept the default and click **Next** to continue.
8. A status window is displayed, indicating the progress of the installation.

9. Once the installation is complete, an Operation Complete message will appear. Click **Finish** to end the installation program.

The utilities have now been successfully installed.

5.2.3 Installing the Storage Manager Client

The Storage Manager Client is used to configure and manage your storage servers. There is also an event monitoring service which is part of the installation. In this section we describe how to install the software.

We recommend that you install the client software onto a dedicated machine that is powered on at all times. You should also consider installing the Storage Manager Client on at least two machines, preferably dedicated management workstations or machines that have access to your storage server using an Ethernet connection. By having two machines configured as management terminals, in the event that something happens on one of these machines, you have quick access to your storage server using the other machine.

Tip: Install the software on your laptop or desktop computer and a dedicated PC or server that is used for monitoring.

Follow these steps to install the client:

1. To install the client software run the downloaded installation package. In our example the file name was SM7client-WIN32-0710G505.exe. The Installation Folder window is displayed.
2. Enter the path to the folder where you wish to decompress the installation package and click **Finish**.
3. If the specified folder does not exist, a window will follow. Click **Yes** and the folder will be created for you.
4. Once the package has been decompressed, a message will inform you that the package has been delivered sucessfully. Click **OK** to continue.
5. Using My Computer or Windows Explorer, locate the folder in which the decompressed package was stored. To begin the installation of the client software, run **`setup.exe`**.
6. An information window will follow, click **OK** to continue.
7. The IBM FAStT Storage Manager 7 Welcome window is displayed. Click **Next** to continue.
8. The Event Monitor window is displayed. If you are installing the software onto your laptop or desktop computer click **No**. If you are installing onto a management workstation or monitoring server, click **Yes**.
9. The Choose Destination Location window appears. Specify where the client software is to be installed or accept the default and click **Next** to continue.
10. A status window is displayed, indicating the progress of the installation.
11. Once the installation is complete, an Operation Complete message will appear. Click **Finish** to end the installation program.

The client software has now been successfully installed.

5.3 Using the Storage Manager Client

Now that the software is installed, the next step is to configure your storage server using the client software. Before attempting to use the Storage Manager Client, we recommend you try to ping both controllers from the machine on which Storage Manager is installed to ensure that basic connectivity is in place. Once you are confident you can access the controllers, you can start to configure them.

Follow these steps when using IBM Storage Manager for the first time:

1. Start the Storage Manager Client from its shortcut located on the Start menu.
2. A window is displayed, asking if you wish to perform an automatic discovery. Click **OK** to continue.
3. The Storage Manager 7 Enterprise Management window is displayed, as shown in Figure 5-8:

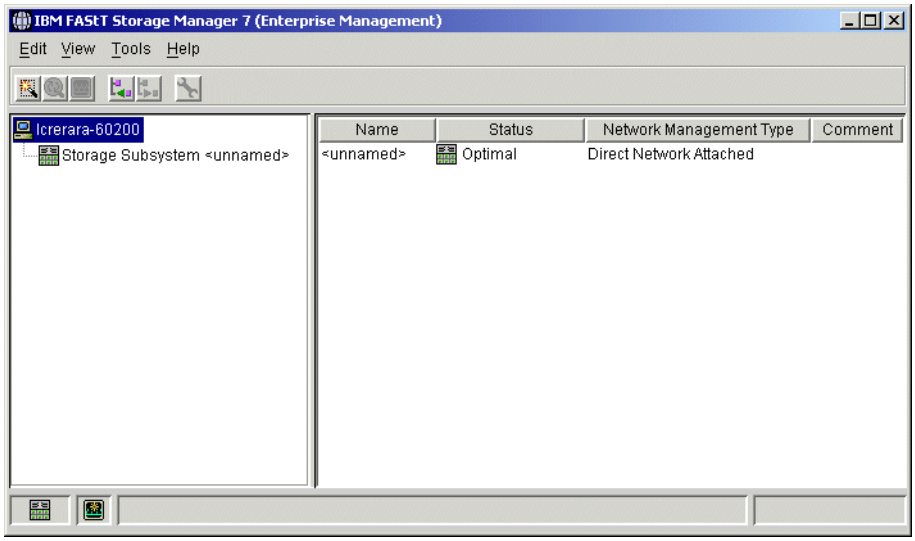

Figure 5-8 IBM Storage Manager 7 (Enterprise Management) window

4. This window displays all discovered storage subsystems. To manage the storage subsystem, right-click **Storage Subsystem <unnamed>**, and select **Manage Device**.
5. The Storage Manager 7 Subsystem Management window is displayed, as in Figure 5-9:

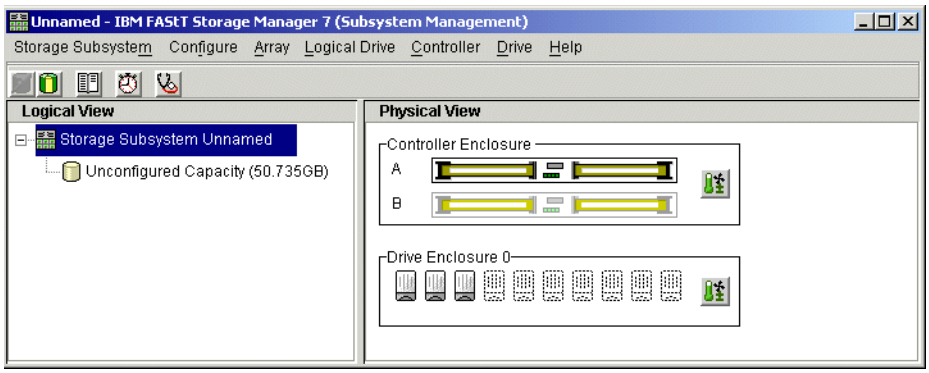

Figure 5-9 Storage Manager 7 Subsystem Management window

This window allow you to manage a specific storage subsystem.

5.3.1 Renaming the storage subsystem

A new storage subsystem is unnamed by default, so the first thing we recommend that you do is to give your storage subsystem a meaningful name.

Tip: Choose a naming standard that can be used with more than one storage server. We recommend that you choose names that are as descriptive as possible, identifying the storage server's location and purpose, for example.

Use the following steps to rename the storage subsystem:

1. Within the Subsystem Management window, right-click **Storage Subsystem Unnamed**, and select **Rename...** (see Figure 5-10):

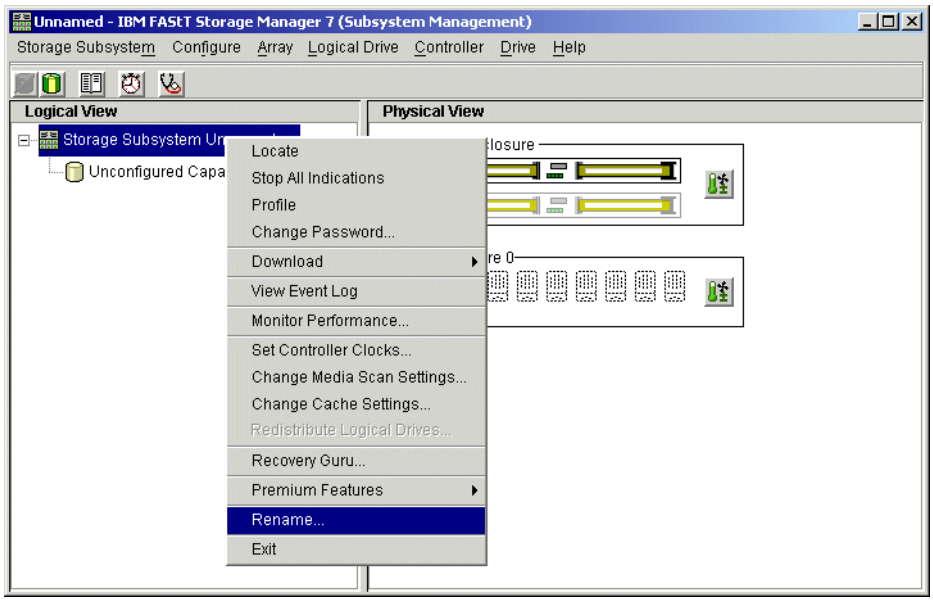

Figure 5-10 Storage Subsystem context menu

2. The window that appears allows you to enter a name for your storage subsystem, as shown in Figure 5-11:

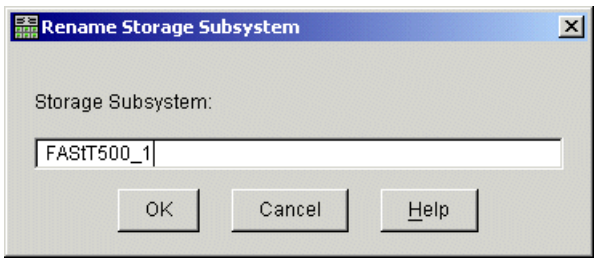

Figure 5-11 Rename Storage Subsystem

3. Type the name for your storage subsystem (we chose `FAStT500_1`), and click **OK**.

The storage subsystem has now been renamed. The new name appears under the subsystem management window as shown in Figure 5-12:

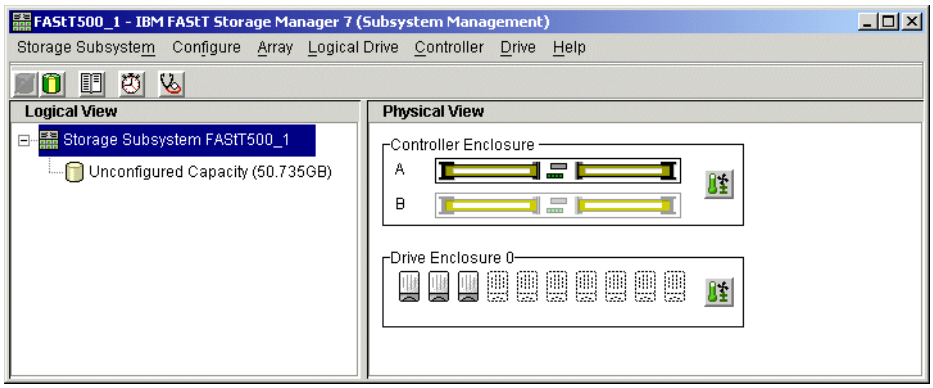

Figure 5-12 FAStT500_1 - IBM Storage Manager 7 (Subsystem Management) window

5.3.2 Changing the storage subsystem password

Initially, the storage subsystem does not have a password. We recommend that you assign a password to prevent unauthorized access to your storage configuration.

Tip: Always change the password of a new storage subsystem. Be wise when choosing a password because, if compromised, an unauthorized person gaining access has rights to destroy your entire configuration. Keep your password secure and change it regularly.

1. Within the Subsystem Management window, right-click **Storage Subsystem FAStT500_1**, and select **Change Password...**
2. The Change Password is displayed as shown in Figure 5-13:

Figure 5-13 Change Password window

3. If you had a previous password you would type it in under Current Password. Since this is a new installation we ignore that and type in our selected password under New Password. We used `password` for simplicity in our example, but you should choose a much better one! You are also required to type in the new password in the Confirm New Password field. Click **OK** to continue.

4. If all fields have been correctly filled in, a `Password Changed Successfully` message is displayed. Click **OK** to complete the procedure.

5.3.3 Setting the controller clocks

Since this is a new installation, we must synchronize the controller clocks to that of the client.

Important: All controller events that are recorded in the Major Event Log (MEL) are time-stamped with the controller's date and time. It is a good idea to synchronize the controller clocks to that of the client (assuming it is correct!) so that the correct date and time are recorded when an event occurs that requires troubleshooting.

Synchonize the clocks using these steps:

1. Within the Subsystem Management window, right-click on **Storage Subsystem FAStT500_1**, and select **Set Controller Clocks...**
2. The Set Controller Clocks window is displayed, as shown in Figure 5-14:

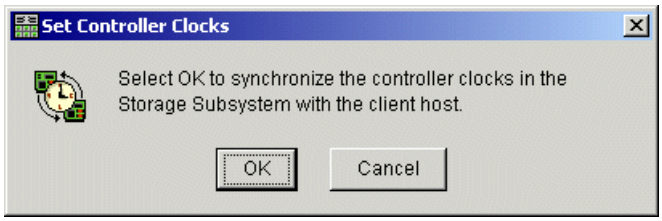

Figure 5-14 Set Controller Clocks window

3. Click **OK** to synchronize the clocks.

5.3.4 Updating the firmware and NVSRAM

Here we discuss the procedures you use to update the FAStT500 controller firmware and Nonvolatile Static Random Access Memory (NVSRAM).

Tip: Always upgrade the firmware and NVSRAM to the latest levels when doing a new installation of a FAStT500 Storage Server. Each new release of firmware holds valuable feature upgrades and bug fixes.

In our example, we upgraded using the firmware package FW_04010207_04010200.dlp and the NVSRAM package NV3552R710NT004.dlp, upgrading from code level 04.01.02.03 to code level 04.01.02.07.

Important: When updating the level of firmware or NVSRAM, remember to check the readme file associated with the update. Any critical information about the update and its specific requirements are included in this file. In our example, the update did not allow for any I/O on either controller.

It is important always to update the firmware before updating the NVSRAM.

Updating the firmware

1. Within the Subsystem Management window, right-click **Storage Subsystem FAStT500_1**, and select **Download->Firmware...**.
2. The Firmware Download window is displayed. Information about the current firmware and NVSRAM levels installed on your storage server is available, as shown in Figure 5-15:

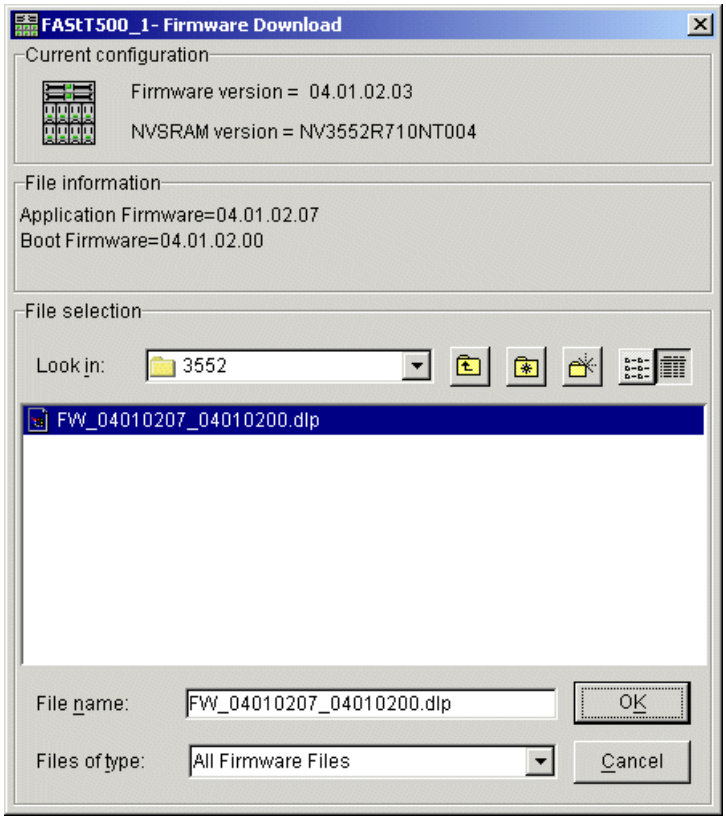

Figure 5-15 Firmware Download window

3. In the lower part of this window, you can select a new firmware package file. After selecting a .dlp file, click **OK** to continue.
4. A pop-up will ask you to confirm that you wish to update the firmware. Click **Yes** to begin the firmware download.
5. A status window displays the status of the firmware download. The process to transfer the firmware and perform the updates may take some time. When the upgrade has been completed for both controllers, the confirmation window shown in Figure 5-16 is displayed:

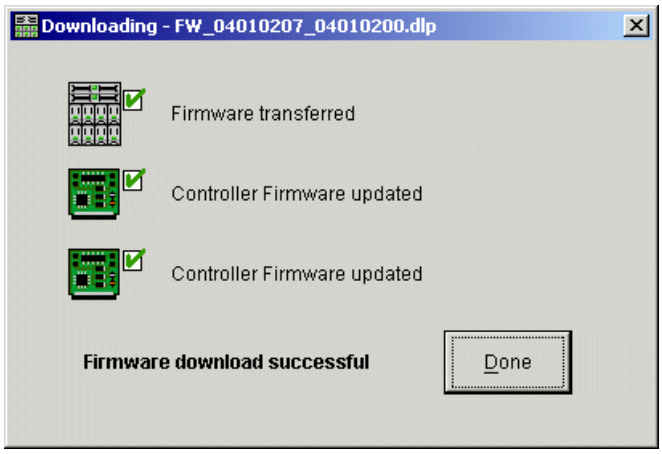

Figure 5-16 Firmware download successful

6. Click **Done** to complete the procedure.

Updating the NVSRAM

Upgrading the NVSRAM is performed in a similar manner:

1. Within the Subsystem Management window, right-click **Storage Subsystem FAStT500_1**, and select **Download->NVSRAM ...**
2. The NVSRAM Download window is displayed. Information about the current firmware and NVSRAM levels installed on your storage server is available in this window.
3. You can select a new NVSRAM package file. After selecting a .dlp file, click **OK** to continue.
4. A pop-up will ask you to confirm that you wish to update the NVSRAM. Click **Yes** to begin the download.
5. A status window displays the status of the NVSRAM download. The process of transferring the NVSRAM and performing the updates may take some time. When the upgrade has been completed, a confirmation window is displayed. Click **Done** to complete the procedure.

After completing both the firmware and NVSRAM updates, you must restart both controllers. You can do this in one of three ways:

- By using the **`sysReboot`** command over the serial connection
- By pushing the reset button on the front panel of the controller
- By restarting the entire storage subsystem

We prefer to use the serial command to perform the reboot, since this allows you to view messages and diagnostic checks as the reboot progresses.

5.3.5 Changing controller modes

You may have noticed that Controller B appears grayed out in the Subsystem Management window, whereas Controller A does not. This is because, although both controllers are online, Controller A is currently active and Controller B is passive.

Controllers are always in one of the following states or modes:

- Online:
  - Active: Can perform IO operation
  - Passive: Cannot perform IO operations, awaiting failure of the other controller
- Offline: Controller is currently down, waiting to go online.

Important: Always have both controllers online to provide redundancy.

To bring Controller B online, follow these steps:

1. Right-click the graphic for **Controller B**, and select **Change Mode->Active** (see Figure 5-17):

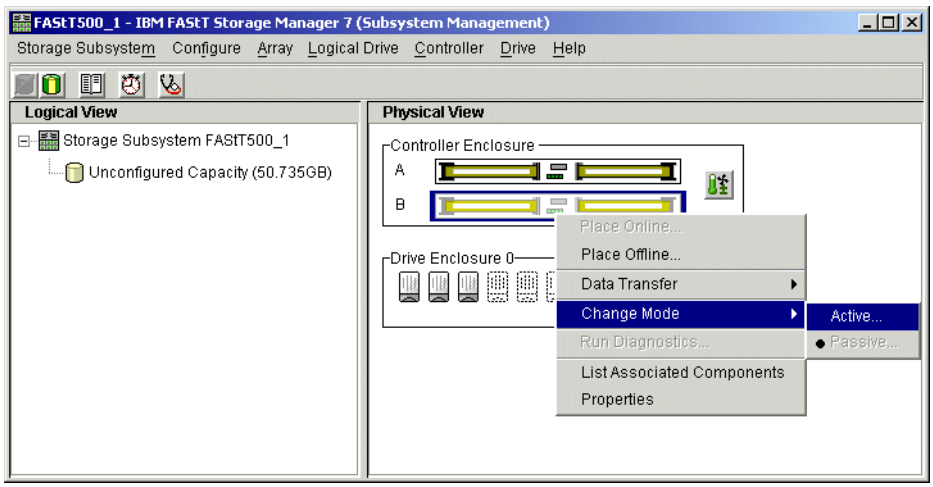

Figure 5-17 Change Mode

2. A window prompting you to confirm your request is displayed. Click **OK**.

The storage subsystem is now running in an active-active configuration.

Tip: For improved performance, we recommend using an active-active configuration rather than an active-passive configuration.

5.3.6 Enabling and disabling Premium Features

A Premium Feature is a storage subsystem feature that is not part of the standard configuration. These features must be purchased as optional extras. In this section we discuss how to enable and disable Premium Features.

Listing Premium Features

To list the current Premium Features, follow these steps:

1. Within the Subsystem Management window, right-click **Storage Subsystem** `FAStT500_1`, and select **Premium Features->List**..
2. The Premium Features window is displayed. This shows the current Premium Features that are enabled, along with the Feature Enable Identifier, as shown in Figure 5-18:

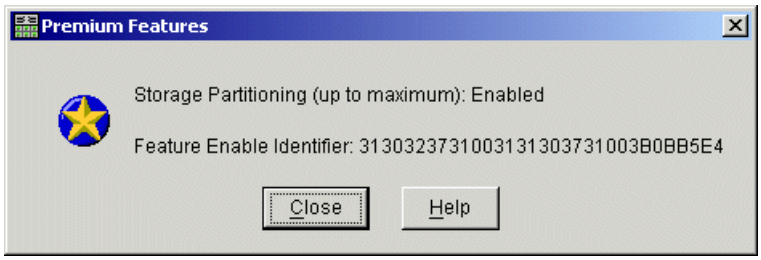

Figure 5-18 Premium Features window

If you wish to enable a Premium Feature that is currently disabled, you have to quote the Feature Enable Identifier code that is listed when requesting a key.

3. Make a note of the relevant Feature Enable Identifier and click **Close**.

Enabling Premium Features

To enable a Premium Feature using a key file, follow these steps:

1. Within the Subsystem Management window, right-click **Storage Subsystem** `FAStT500_1`, and select **Premium Features-> Enable**.
2. The Select Feature Key File is displayed (see Figure 5-19):

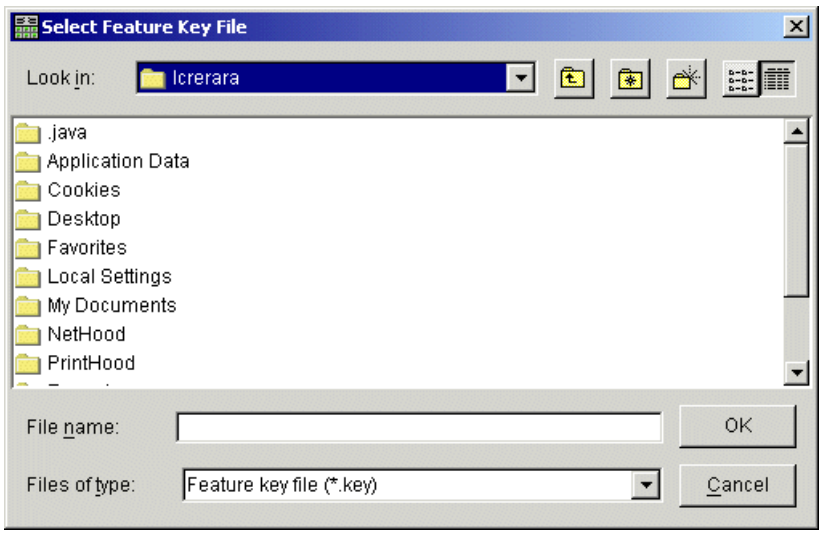

Figure 5-19 Select Feature Key File window

3. Here we specify the .key file that was supplied to us. Select the file and click **OK**.
4. The Confirm Enable Feature Key window is displayed. Click **Yes** to enable the Premium Feature.

Disabling Premium Features

To disable a Premium Feature, follow these steps:

1. Within the Subsystem Management window, right-click **Storage Subsystem** `FAStT500_1`, and select **Premium Features->Disable**.
2. The Disable Premium Feature window is displayed (see Figure 5-20):

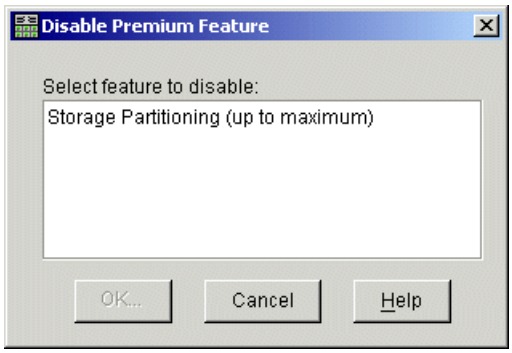

Figure 5-20 Disable Premium Feature window

3. Select the Premium Feature you wish to disable, and click **OK**.
4. A confirmation window is displayed. Click **Yes** to disable the Premium Feature.
5. A Working message will follow, which closes once the Premium Feature has been disabled.

5.3.7 Creating arrays and logical drives

The next step in getting your SAN up and running is to create an array and a logical drive. First, let us define those terms:

- Array

  An arrangement of disk drives that you have grouped together and assigned a specific RAID level.

- RAID

  The ability to stripe or mirror data across multiple disks, to form larger and, perhaps, redundant disk configurations. A number of different RAID levels are available, which meet varying requirements for redundancy, performance and capacity.

- Logical Drive

  A logical drive is created within the bounds of an array and is seen as a physical disk by host operating systems. Logical drives may be formed from an entire array or just a portion of it. A logical unit number (LUN) is assigned to a logical drive.

Follow these steps to create an array and logical drive:

1. Within the Subsystem Management window, right-click **Unconfigured Capacity**, and select **Create Array/Logical Drive...** (see Figure 5-21):

5. The Create Logical Drive - Specify Logical Drive Parameters window is displayed (see Figure 5-23):

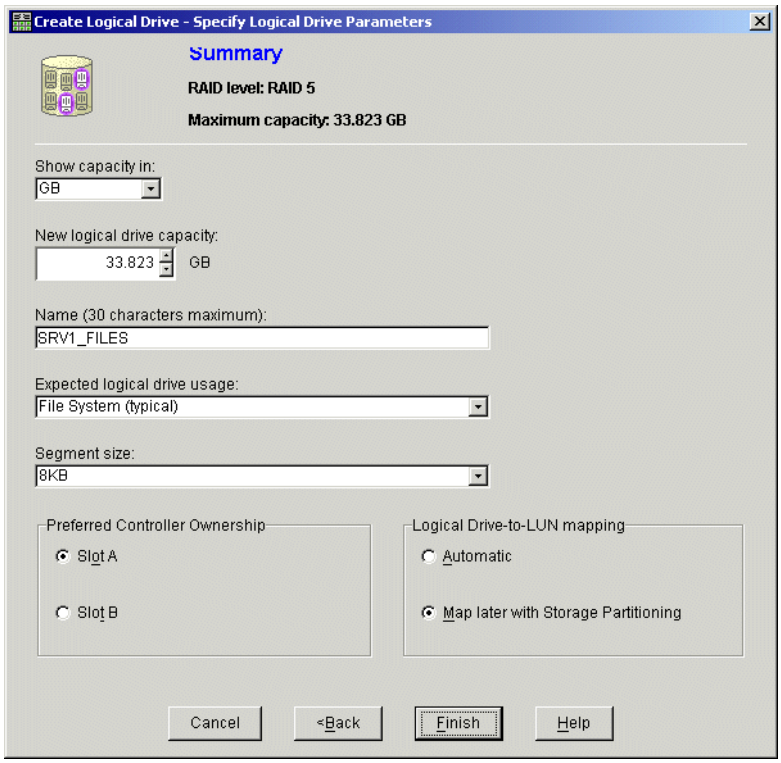

Figure 5-23 Create Logical Drive - Specify Logical Drive Parameters window

6. This window lets you specify a number of different parameters for the logical drive you are about to create.

   In our example we chose to use the full capacity of the array, named our logical drive SRV1_FILES, set Expected logical drive usage to File System (typical), set segment size to 8K, set Preferred Controller Ownership to Slot A and set the Logical Drive to LUN mapping to Map later with Storage Partitioning. When you have made all of your selections, click **Finish**.

The array and logical drive are now created. Initialization commences, which takes some time. However, the logical drive is now available to be assigned to a server with storage partitioning.

5.3.8 Storage partitioning

Now that we have created our logical drive, we move on to discuss storage partitioning, which is also known as LUN masking.

Storage partitioning is interesting because it allows your SAN to operate in a heterogeneous environment. That is, it allows you to connect multiple servers, running different operating systems, to your storage server. By defining hosts and groups, you are able to hide LUNs or assign them to particular servers. In this way, multiple servers connected to your SAN can access different disk resources on a single storage server.

Follow these steps to configure storage partitioning:

1. From the Subsystem Management window menu, select **Configure->Storage Partitioning...**.
2. The FAStT500_1 - Mappings window is displayed, which allows you to perform storage partitioning (see Figure 5-24):

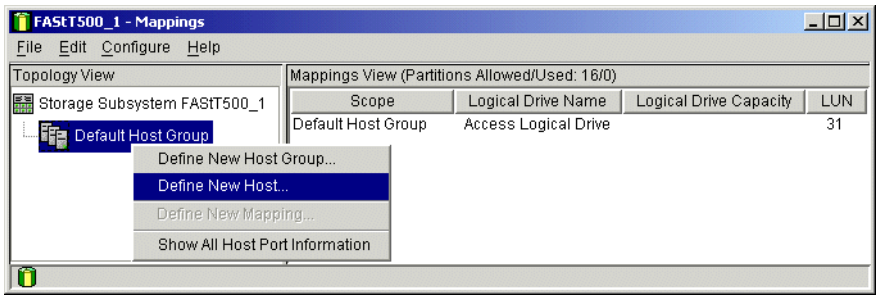

Figure 5-24 Mappings window

3. Right-click **Default Host Group**, and select **Define New Host...**.
4. The Define New Host window is displayed. Enter the name of your server (we used SRV1) and click **Add** to continue (see Figure 5-25):

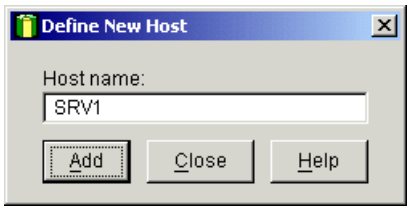

Figure 5-25 Define New Host window

5. The new host is added to the configuration as shown in Figure 5-26:

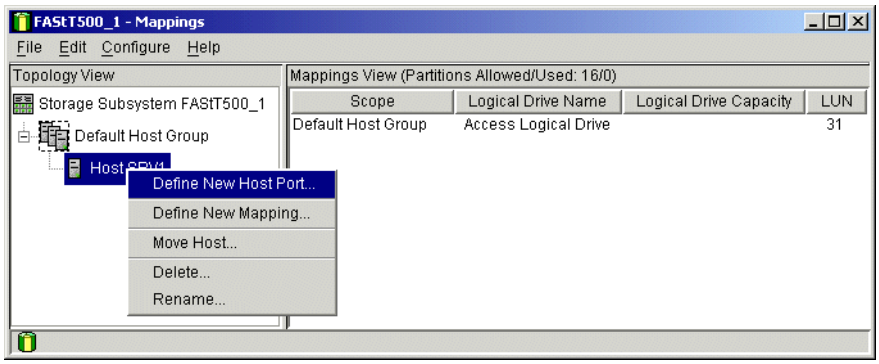

Figure 5-26 Defining a new host port

6. Right-click SRV1, and select **Define New Host Port...**.
7. The Define New Host Port window will open, as shown in Figure 5-27:

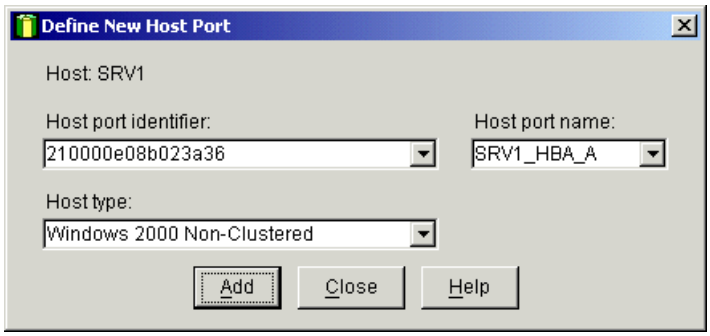

Figure 5-27 Define New Host Port window

8. Select the host port identifier or WWN of an HBA in your server, and the host type (which is the host operating system), and specify the host port name. Click **Add** to continue.
9. Repeat steps 6 and 7 if your server contains more than one host bus adapter.
10. Right-click **SRV1**, and select **Define New Mapping...**.
11. The Define New Mapping window is displayed, as shown in Figure 5-28:

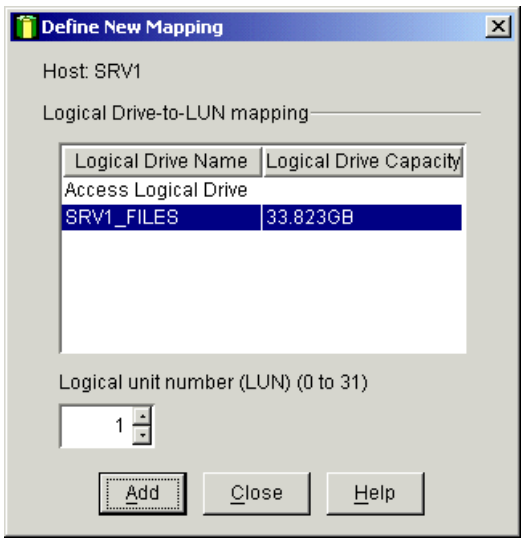

Figure 5-28 Define New Mapping window

12., Select the logical drive and specify the LUN number. Click **Add** to map the drive to SRV1.
13. Close the FAStT500_1 - Mappings window.

5.3.9 Changing media scan settings

Media scan is a background process that can be enabled for logical drives in the storage subsystem. The advantage to enabling media scan is that it provides error detection, helping to detect errors before normal drive operation is interrupted. We recommend that you enable media scan on all logical drives.

We must enable media scan on the storage subsystem and we must set the interval, which is how often we want media scan to run.

Perform the following to enable media scan:

1. Within the Subsystem Management window, right-click **Storage Subsystem FAStT500_1**, and select **Change Media Scan Settings...**.
2. The Change Media Scan Settings window is displayed, as shown in Figure 5-29:

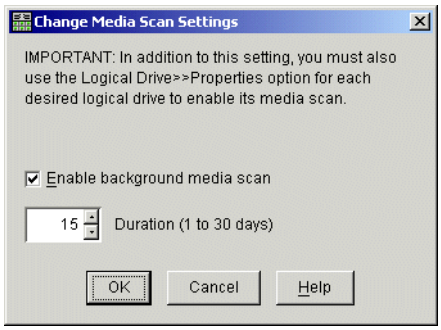

Figure 5-29 Change Media Scan Settings Dialog Box

3. Check the Enable background media scan check box.
4. Enter a value for Duration, which determines how often the media scan runs. We left the duration value at a 15-day interval.
5. Click **OK** to commit the changes.

Media scan has now been enabled on the storage subsystem. Note that logical drives must also have media scan enabled individually for the scan to run on a particular logical drive. We show how this is done in the next section.

5.3.10 Enabling media scan for a logical drive

Here we show you how to enable media scan for a specific logical drive. This follows on from the previous section on media scan settings for the whole subsystem.

To enable media scan on a logical drive perform the following:

1. Right-click the logical drive, in our example, **Logical Drive SRV1_FILES**, and select **Properties**.
2. The Logical Drive Properties window is displayed, as shown in Figure 5-30:

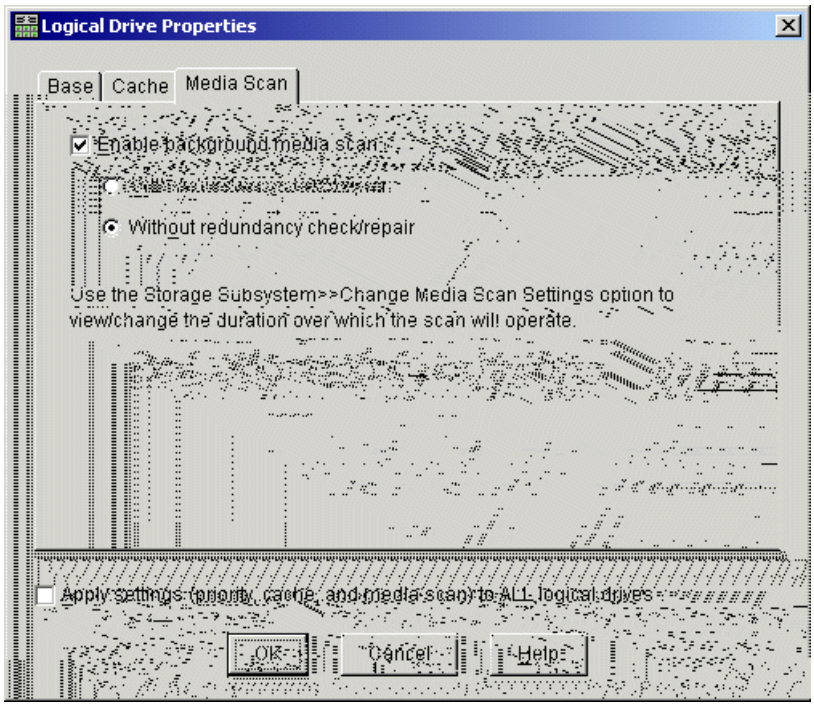

Figure 5-30 Logical Drive Properties

3. Click the **Media Scan** tab.
4. Check **Enable Background Media Scan**. Enable the radio button **Without redundancy check/repair**.
5. Click **OK** to enable the media scan settings on the logical drive.

Remember to enable media scan on all logical drives when you create them.

We do not discuss the remaining logical drive properties. For in-depth information on these settings, please refer to the Storage Manager online help or product documentation.

5.3.11 Saving the configuration

It is a good practice to save your configuration. This option outputs a script file that can be used to recreate your entire configuration as it appeared at that moment in time. We recommend saving the configuration after making any configuration changes to your storage subsystem.

Do not save the configuration while changes are still being executed, such as RAID changes.

Tip: Implement a change control procedure for your storage subsystem, including saving the configuration and profile every time a change is made. This will be invaluable in a disaster recovery situation. It can also be used to help you to create a replica storage subsystem at an alternate site.

To save the storage subsystem configuration, perform the following steps:

1. From the Subsystem Management window menu, select **Configure->Save Configuration...**.

2. An information window is displayed, informing you that you should not save the configuration if there is a change operation still running. If not, click **Yes** to continue.
3. The Save Configuration window is displayed. Specify the name and location of the configuration file that you wish to create and click **Save**.

The storage subsystem configuration is saved to the specified location.

5.3.12 Saving the profile

The Storage Subsystem Profile is a description of the entire subsystem configuration, including logical and physical drive information. It is important to save the profile after making changes, since it gives you the ability to check the configuration of your storage server at a moment in time, forming part of a historical record of changes to the configuration. It may also be asked for by IBM Support should you need assistance in troubleshooting any problems.

To save the storage subsystem profile, perform the following steps:

1. Within the Subsystem Management window, right-click **Storage Subsystem FAStT500_1**, and select **Profile**.
2. The Storage Subsystem Profile window is displayed, as shown in Figure 5-31:

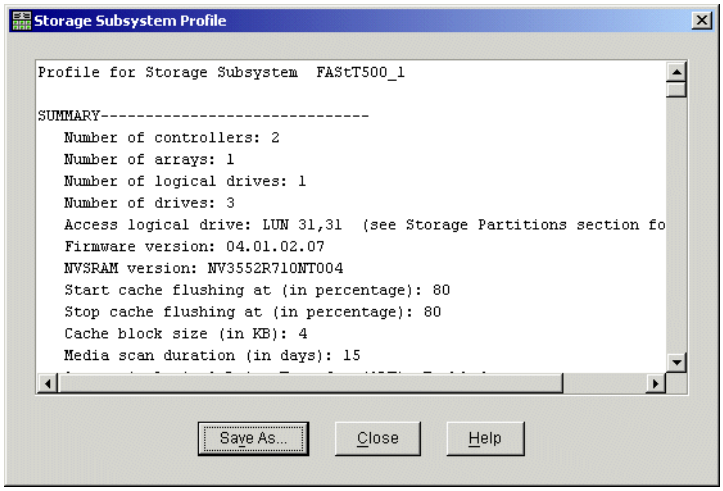

Figure 5-31 Storage Subsystem Profile

3. The current profile of your storage subsystem is displayed as readable text. To save the profile to a file, click **Save As...**.
4. The Save Profile is displayed. Specify the name and location of the file that you wish to create and click **Save**.

The storage subsystem profile is saved to the specified location.

5.3.13 Using the Storage Manager utilities

The **SM7devices.bat** command is used to display all devices available to the server using the SAN and FAStT Storage Server. This is particularly useful when troubleshooting and you need to know which LUNs are available to a server.

Use the following steps to execute the **SM7devices.bat** command:

1. Open a Command Prompt by clicking the icon in the Start menu or by running **cmd.exe**.

2. Change to the directory in which you have installed the utilities. For example:

   ```
   cd C:\Program Files\SM7util
   ```

3. Execute the command **SM7devices.bat**. The output of this command shows which LUNs are assigned and available to the server. If you look carefully you will see that the storage server, logical drive name, LUN and WWN associated with the device are also displayed. Sample output is shown in Figure 5-32:

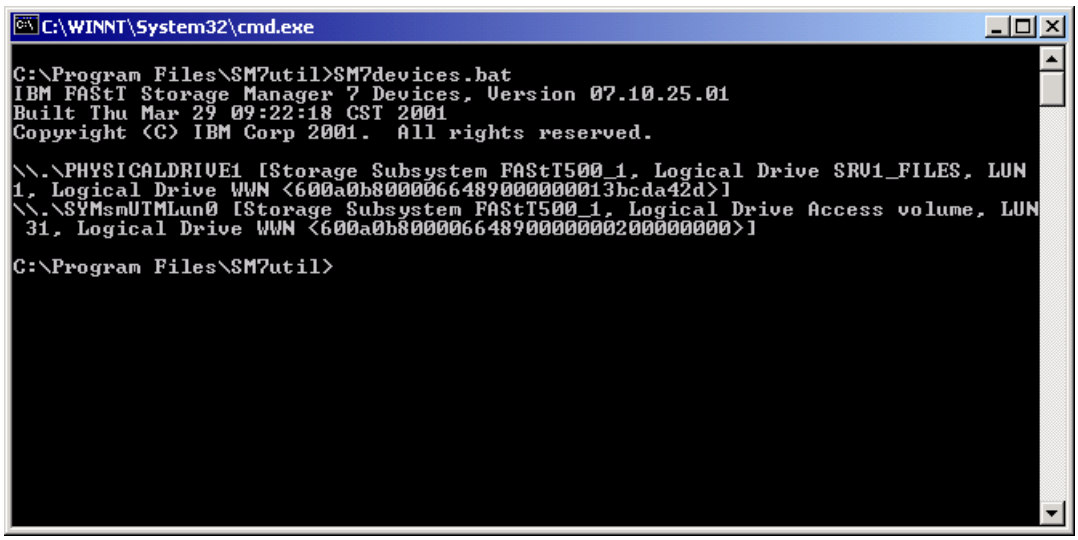

Figure 5-32 SM7devices.bat window

The **hot_add.exe** command is used to initiate detection of LUN changes at the operating system level. This is also known as hot-adding disks to the operating system. This command can be used to detect new LUNs created and assigned to a server, without the requirement for a server reboot.

Use the following steps to execute the **hot_add.exe** command:

1. Open a Command Prompt by clicking the icon in the Start menu or by running **cmd.exe**.
2. Change to the directory where you have installed the utilities. Execute the command **hot_add.exe**. The operating system should now detect if any new LUNs have been added.
3. Using Windows 2000, we can check that the new LUN has been detected in Computer Manager. Right-click **My Computer** on the desktop and select **Manage**.
4. The Computer Management window is displayed. Right-click **Disk Management**, and select **Rescan Disks**.
5. The Windows 2000 Write Signature and Upgrade Disk Wizard is invoked. Follow the normal Windows process and the new drive will be visible in the disk manager window. Use the standard Windows procedures for creating a new partition.

Tip: Use **hot_add.exe** to reduce downtime when creating and assigning new LUNs to a server.

6

FAStT management

This chapter discusses FAStT Storage Manager, the software that provides you with sophisticated management capabilities for organizing your SAN. We describe how to perform a number of useful functions, including:

- Configuring alerts to be sent using e-mail or Simple Network Management Protocol (SNMP).
- Using the Recovery Guru to identify problems and execute recovery procedures.
- Using the Major Event Log (MEL) in conjunction with the Recovery Guru.
- Using performance monitor to collect performance data.

We also provide a brief introduction to the script editor, which can be used to automate configuration of the storage subsystem.

6.1 Configuring alerts

One of the key elements of managing a storage subsystem is detecting problems that may arise. To avoid having to nominate someone to perform the onerous task of continually monitoring the subsystem, it is possible to configure automatic alerting so that you are notified if any out-of-the-ordinary situations arise.

Alerting allows you to configure your storage subsystem to send a message to nominated parties, either using e-mail or SNMP, when something goes wrong. No time is wasted checking the health of the storage subsystem. This leaves more time to be spent on performance tuning and disk management.

Follow these steps to configure alerting for the storage subsystem:

1. In the Enterprise Management window menu, select **Edit->Configure Mail Server...** (see Figure 6-1):

© Copyright IBM Corp. 2002

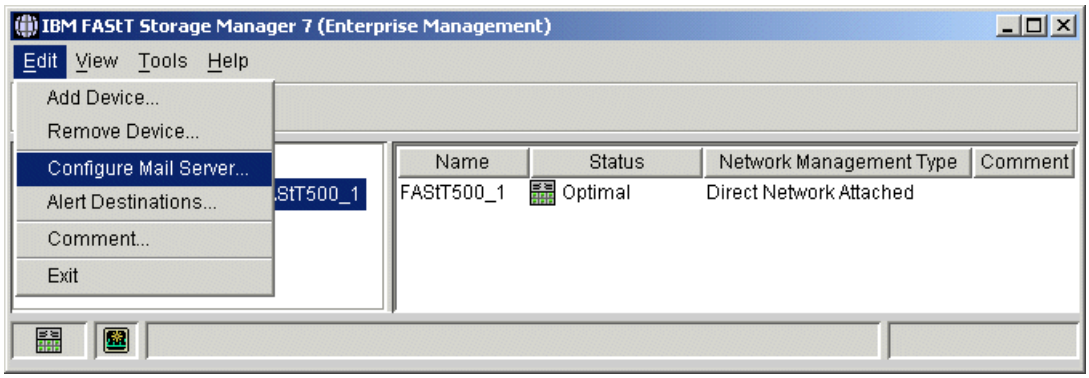

Figure 6-1 Enterprise Management window

2. The Configure Mail Server window appears, as shown in Figure 6-2:

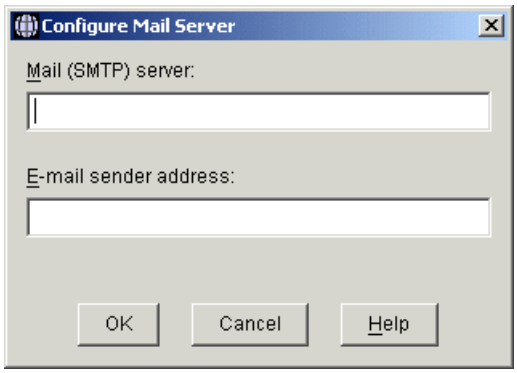

Figure 6-2 Configure Mail Server window

3. In the Mail (SMTP) server text box, enter the IP address or DNS name of your outgoing mail server. In the remaining box, enter the e-mail address that will appear as the sender of the alerting e-mail, such as `FAStT@itso.ibm.com`. After completing the text boxes, click **OK**.
4. Select **Edit->Alert Destinations...** to display the Add/Edit Alerts window (see Figure 6-3):

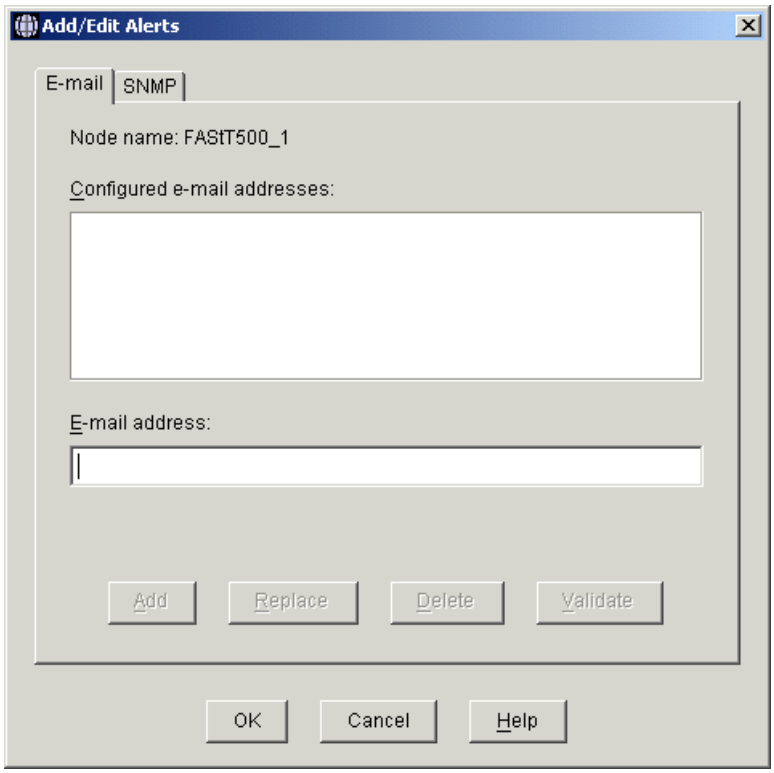

Figure 6-3 Add/Edit Alerts E-mail tab

5. Here, you specify the recipients of the e-mail alerts. In the E-mail address field, type in an e-mail address to which you wish to send the alerts, and click **Add** to add the address to the Configured e-mail address list.
6. Repeat step 5 until you have added all of the addresses to which you wish to send e-mail alerts.
7. If you do not require SNMP alert support, click **OK** to close the window.
8. If you require SNMP alert support, click the **SNMP** tab.
9. The SNMP tab is displayed (see Figure 6-4):

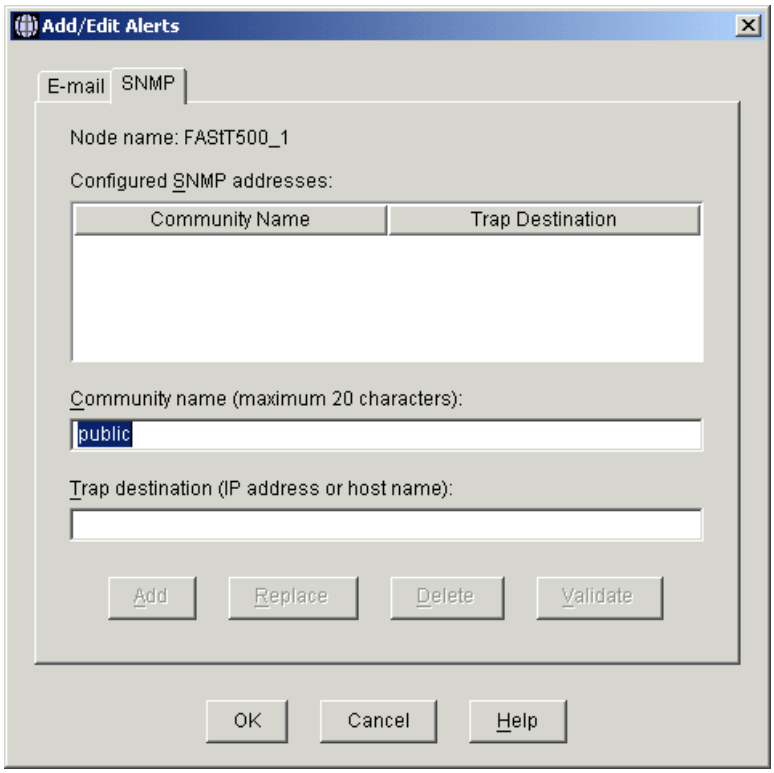

Figure 6-4 Add/Edit Alerts SNMP tab

10. In the Community name text box, enter the community name. In the Trap destination text box, enter the IP address or DNS name of the system to which you wish to send the SNMP trap. Click **Add** to add the SNMP address to the Configured SNMP addresses list.

    For more information on SNMP please refer to the Simple Network Management Protocol (SNMP), RFC 1157 at

    `http://www.rfc-editor.org`

11. Repeat step 10 if you require more SNMP address entries.
12. Click **OK** to close the Add/Edit Alerts window.

If a critical event occurs, such as a drive failure, an alert will be e-mailed to the configured e-mail addresses, and an SNMP trap will be forwarded to the specified SNMP management stations.

6.2 Recovery Guru

The Recovery Guru is a tool that helps identify storage subsystem problems and recommends a procedure to follow when attempting to fix these problems. This is useful if you have never before been exposed to storage subsystem problems, and do not have experience in correcting them.

Follow these steps to use the Recovery Guru:

1. To start the Recovery Guru from within the Subsystem Management window, right-click **FAStT500_1 Storage Subsystem**, and select **Recovery Guru...**.
2. If no problems are discovered within the storage subsystem, a message is displayed to this effect. Click **OK** and stop.

3. If there is a problem within the storage subsystem, as in our example, the Recovery Guru window will be displayed, as shown in Figure 6-5:

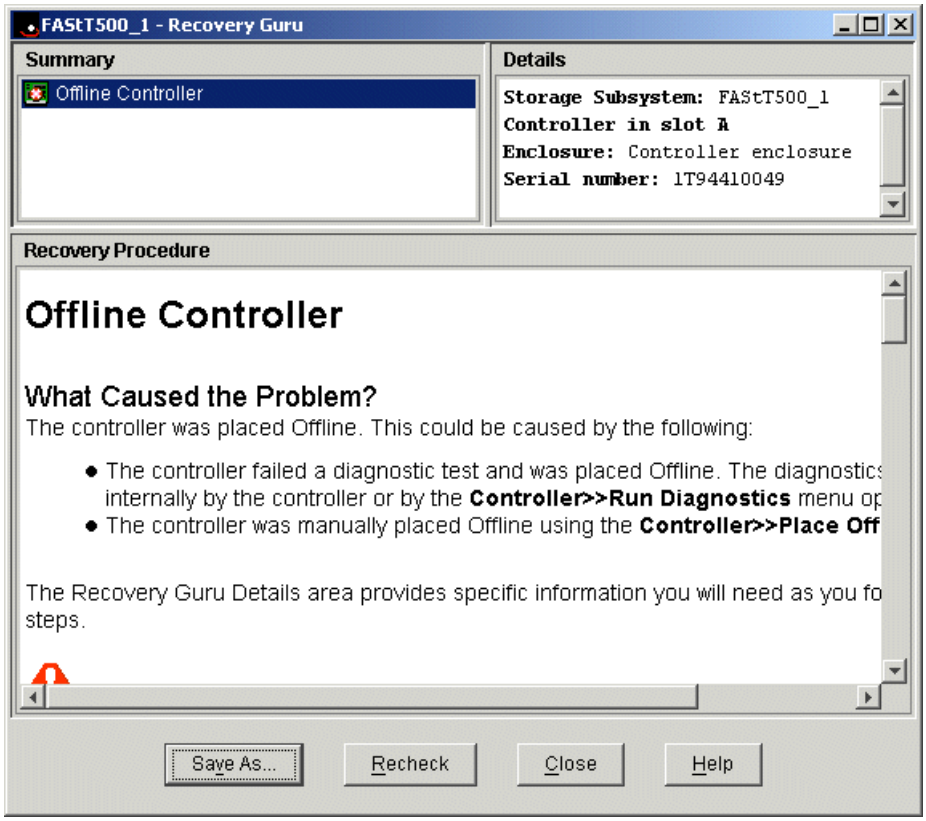

Figure 6-5 Recovery Guru window

4. As you can see, an error condition has been detected and the Recovery Guru displays a summary of faults and details of the selected fault. The Recovery Guru also suggests a recommended recovery procedure to correct this fault.
5. To save the problem to a file, click **Save As...** and follow the on-screen instructions.
6. After taking corrective actions, you can check to see if the problem still exists by clicking **Recheck**..
7. To close the Recovery Guru, click **Close**.

For more details on using the Recovery Guru please refer to the Storage Manager online help system or product documentation.

6.3 The Major Event Log

When a problem arises and you need to understand the cause or reason for the problem, the FAStT subsystem's Major Event Log (MEL) comes into play. The MEL resides on the FAStT Storage Server and is useful when troubleshooting problems or understanding what the storage subsystem is doing. It also allows you to track events and alerts over a specific period. The Recovery Guru should be the first resource you utilize when trying to solve storage subsystem problems, but the MEL can be helpful when used in conjunction with it.

The Storage Manager online help system can be helpful when you have to troubleshoot unknown events.

To view the storage subsystem event log perform the following steps:

1. Within the Subsystem Management window, right-click **Storage Subsystem FAStT500_1**, and select **View Event Log** (see Figure 6-6):

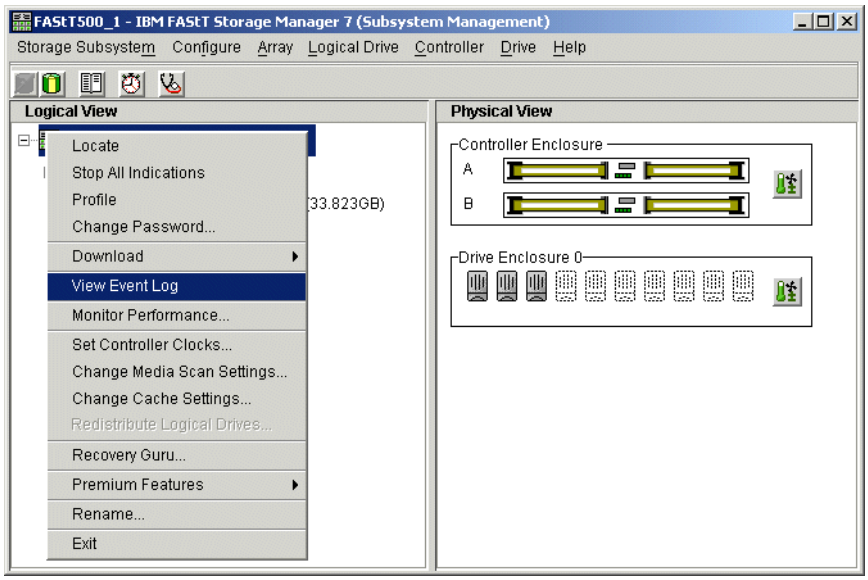

Figure 6-6 Subsystem Management window

2. The FAStT500_1 - Event Log window is displayed, as shown in Figure 6-7:

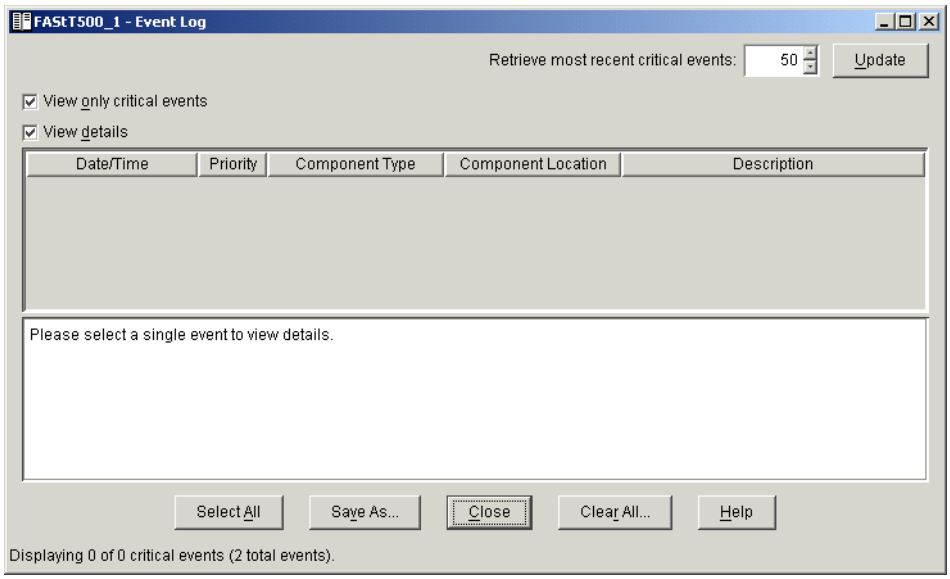

Figure 6-7 Event Log window

3. This window gives you access to a number of event management features. From within the Event Log window you can perform the following tasks:
   a. Set the number of events you wish to retrieve:
      i. The total number of events available in the entire event log is indicated at the bottom left-hand corner of the window. Set the **Retrieve most recent critical** events box to the number you wish to view.

ii. Click **Update** to display the events.

b. Filter between critical and all events:

To see only critical events, check **View only critical events**. To see all events, clear the **View only critical events** check box.

c. View the details of events:

Check **View Details** and the details of a selected event in the upper pane are displayed in the lower pane.

d. Select events and delete them:

i. Select the desired events by clicking them.

ii. To delete the events, press the Del key on the keyboard.

e. Select events and save them to a file:

i. Select the desired events by clicking them.

ii. Click **Save As...**.

iii. The Save Events window is displayed. Enter the path and file name for the file in which you wish to save the events.

iv. Click **Save** to save the events to an ASCII text file.

f. Clear all events:

i. To clear all events, click **Clear All...**.

ii. A confirmation window will follow. Click **Yes** to clear all events.

> **Note:** We do not recommend clearing the event log without saving it first. Not having the information available in the event log can make it difficult to troubleshoot a storage subsystem problem. We recommend saving all events to a file before clearing the MEL.

4. To close the FAStT500_1 - Event Log window, click **Close**.

6.4 Performance monitoring

Having the ability to monitor the performance of your storage subsystem is crucial to being able to utilize its resources to their full potential. The performance monitor allows you to see what types of I/O operations are being performed by the storage subsystem. With this information you are able to identify when disk resources are either under- or overutilized, perhaps causing performance problems. It is a useful tool when undertaking performance tuning of both your system and overall environment.

A detailed performance tuning discussion is outside the scope of this Redpaper. We recommend *Tuning Netfinity Servers for Performance - Getting the most out of Windows 2000 and Windows NT 4.0*, SG24-5287 if you require further guidance. It can be obtained from:

`http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg245287.html`

To start Performance Monitor, follow these steps:

1. Within the Subsystem Management window, right-click **Storage Subsystem `FAStT500_1`**, and select **Monitor Performance...**.
2. The FAStT500_1 - Performance Monitor window is displayed (see Figure 6-8):

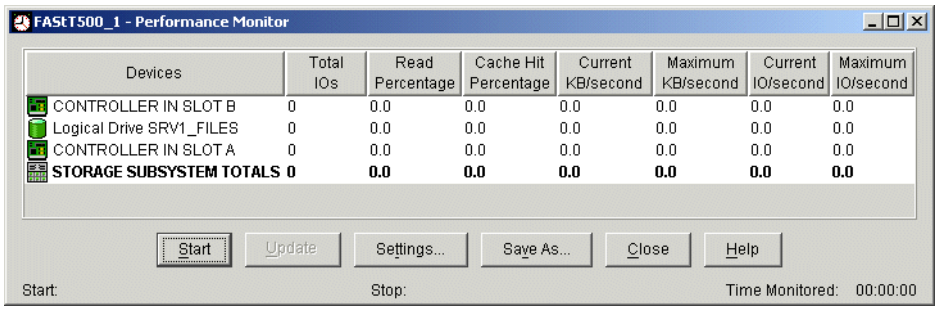

Figure 6-8 Performance Monitor window

3. From within the Performance Monitor window we can perform the following tasks:
   - Click **Start** to commence capturing performance data.
   - Click **Stop** to cease capturing performance data.
   - Click **Update** to perform a one-time capture of performance data.
   - To determine when and which data is captured, click **Settings...**. The Performance Monitor Settings window is displayed (see Figure 6-9):

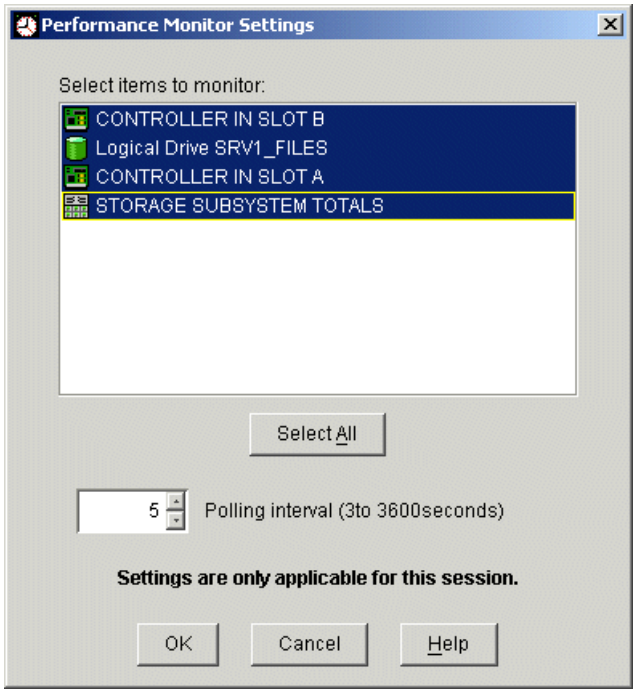

Figure 6-9 Performance Monitor Settings window

   In this window, you can set the polling interval and select which items you wish to monitor. After making changes, click **OK** to put them into effect.

g. Click **Save As...** to be presented with the Save Performance window. This allows you to save performance data in either an ASCII or Comma Delimited file format. Make your selection, enter the path and filename (the extension is .perf) and click **Save**.

h. Click **Close** to close the Performance Monitor window.

i. Click **Help** for detailed information about the columns containing performance data.

Note: The Performance Monitor window must be open to collect performance data.

6.5 The Script Editor

Use of the FAStT Script Editor is outside the scope of this paper. However, we felt that it was important to mention that FAStT Storage Manager has a built-in scripting language, since it can be used to create scripts to automate the creation and configuration of a storage subsystem.

Important: Scripting is a very powerful tool when used correctly. It can also cause severe problems if used incorrectly. Please carefully read all relevant documentation before attempting to use the Script Editor.

To open the Script Editor from within the Enterprise Management window, select **Tools->Script Editor**. Figure 6-10, shows a basic script that has been executed in the Script Editor window.

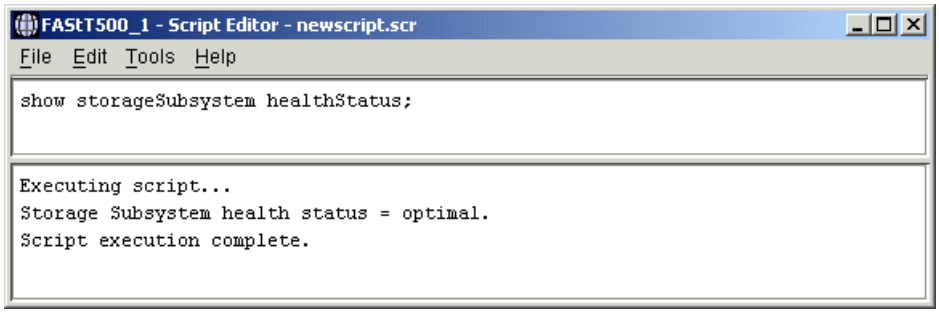

Figure 6-10 Script Editor window

For more information on Script Editor and scripting within FAStT Storage Manager, please refer to the product documentation and online help system.

7

Fabric expansion

In this chapter we demonstrate the ease with which you can expand your SAN fabric by cascading (interconnecting) SAN switches. Scalability is a desirable attribute for a SAN to possess, and IBM SAN Storage Servers offer significant scalability. If you have added devices to your SAN and the number of available ports through which you can connect to it is getting perilously low, cascading switches can solve your problem. Cascading switches is also a way to increase the level of path redundancy within the fabric and, if desired, to increase the distance between server and FAStT Storage Server.

7.1 Cascading switches

Cascading is a cost-effective and easy way to increase port counts in your SAN fabric. Linking SAN switches together provide high resiliency and fault tolerance in the SAN by eliminating a single point of failure.

Adding an inter-switch links (ISL) automatically initiates updates of routing and zoning information across all ISLs. This is an auto-discovery feature of the Fabric OS. It does this by using a Distributed Name Server. The name server is fully distributed to each switch, therefore ensuring no single point of failure. When any end node devices want to communicate with other nodes in the fabric, any switch can provide the necessary connection information by means of the Distributed Name Server.

The Fabric OS uses dynamic route recovery, so, even if a switch fails, connectivity is maintained.

7.2 Configuration tasks

In the configuration we built during the writing of this paper, we started with a single 2109 S16 switch and then expanded our fabric by adding a second S2109 S16, and then a 2109 S08 switch.

Here are the steps we took. Remember we already have one 2109 S16, and we are adding a second of these switches:

© Copyright IBM Corp. 2002

1. Attach a serial cable to the new 2109 S16 switch. For the moment, do not attach any fiber cables or devices.
2. Using a terminal emulator connected to the serial port, set the switch IP address as shown in 4.1, “Assigning an IP address to a Fibre Channel switch” on page 29.
3. Now connect to the switch using the Ethernet connection and log in to the switch using your Web browser.
4. In the Switch Admin window for the new switch, enter a switch name and domain ID. Each switch connected to the fabric must have a unique domain ID (see Figure 7-1):

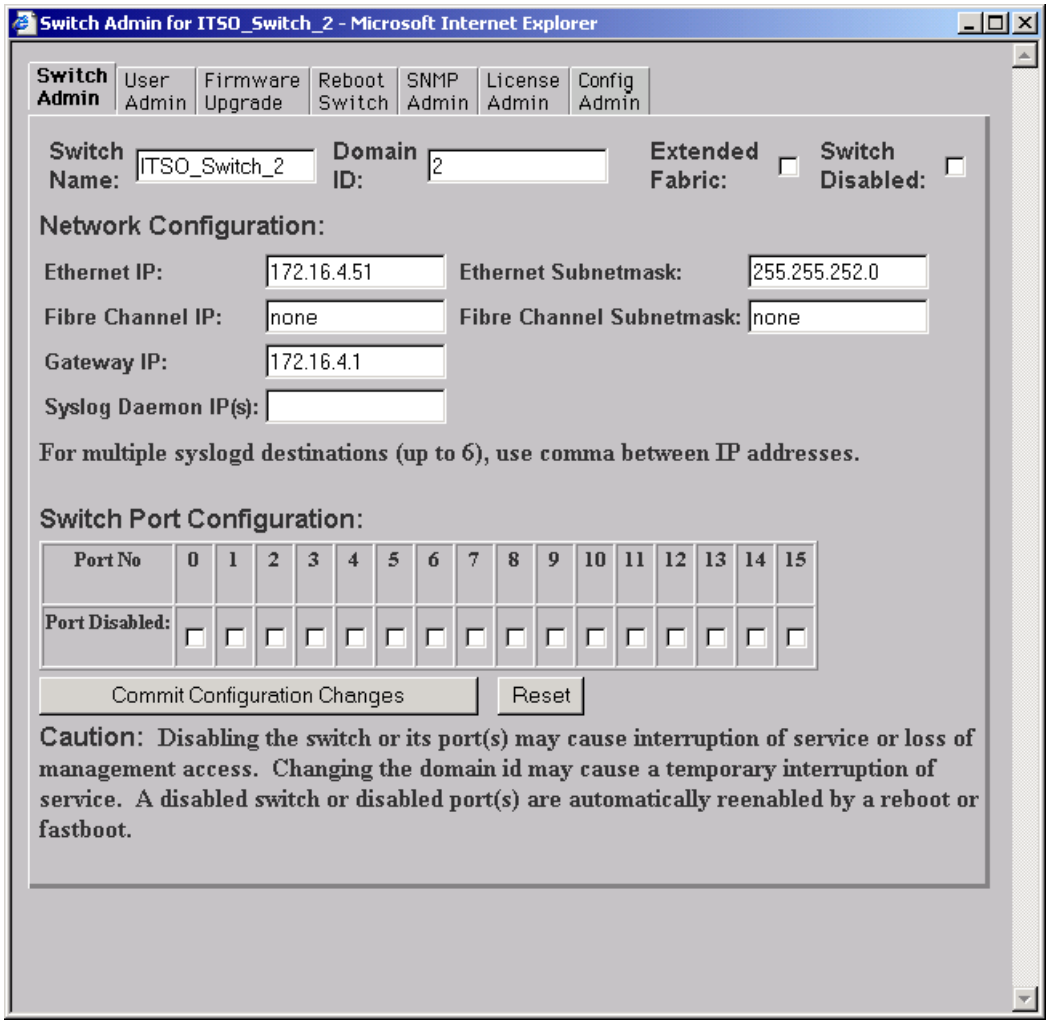

Figure 7-1 Second switch configuration

5. Connect a port on the new switch to one on the existing switch by running a fiber cable between them. This creates an inter-switch link (ISL). There will be a lot of switch activity at this point as the Fabric OS updates the Name Server Table and distributes all zone and port information thoughout the SAN.

   Because we used software zoning, we could easily move some of the devices across to ITSO_switch_2. We did so and still had disk access during this time, since we moved only one cable at a time.
6. Using your Web browser, access the IP address of switch 1. You should now see both switches, as shown by Figure 7-2:

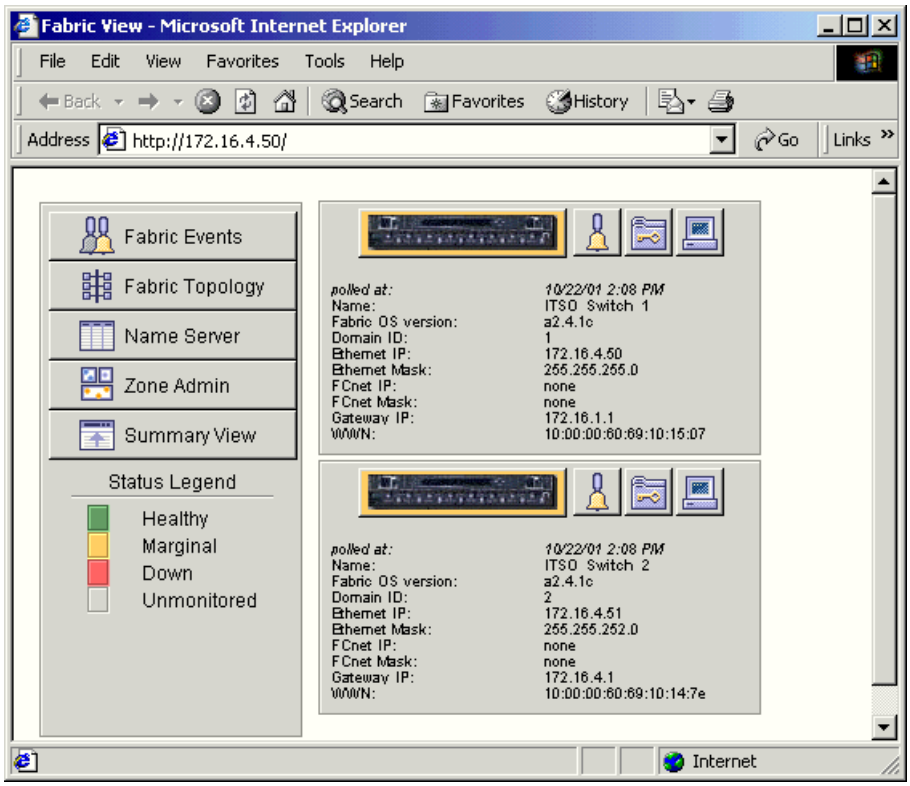

Figure 7-2 The second switch is now visible in the fabric

7. We repeated steps 1 through 6 for a third switch, a 2019 S08. The third switch can be seen in Figure 7-3:

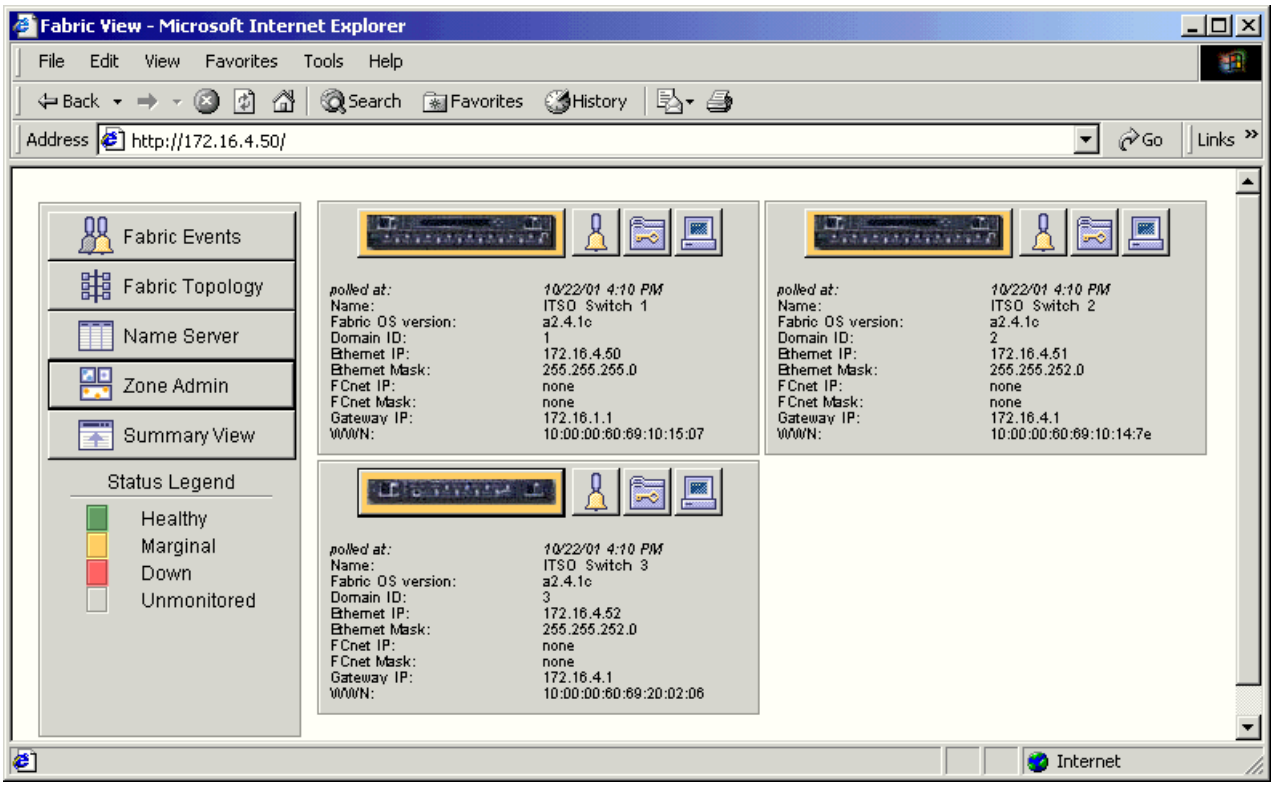

Figure 7-3 Three switches in the fabric

We then moved some existing connections to this switch, forming a *meshed fabric*, where connections to devices are distributed across the switches. The Name Server Table (shown in Figure 7-4) shows the ports used on each switch in our configuration, along with the devices attached to them. You can see domain IDs 1, 2, and 3 which correspond to the switches, as already discussed.

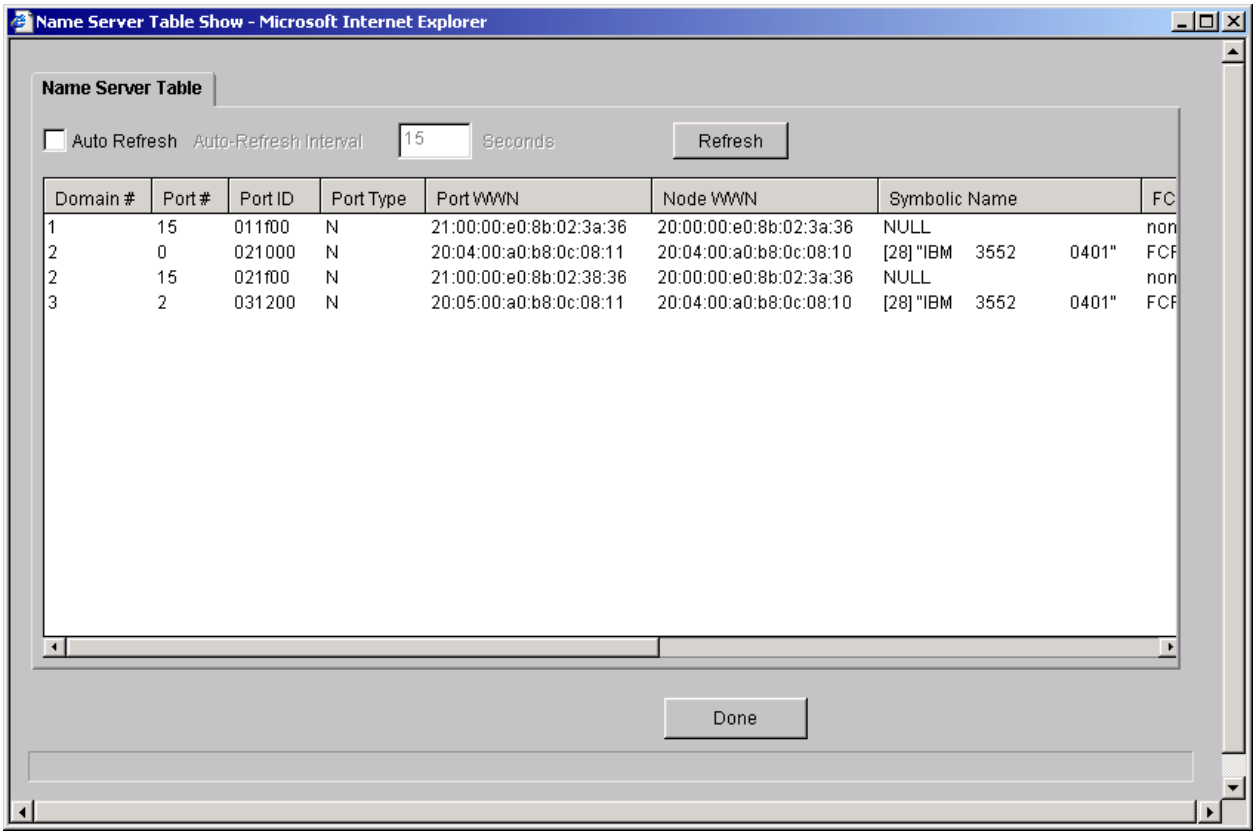

Figure 7-4 SAN Name Server Table

Figure 7-5 shows the topology of the SAN configuration after cascading the switches.

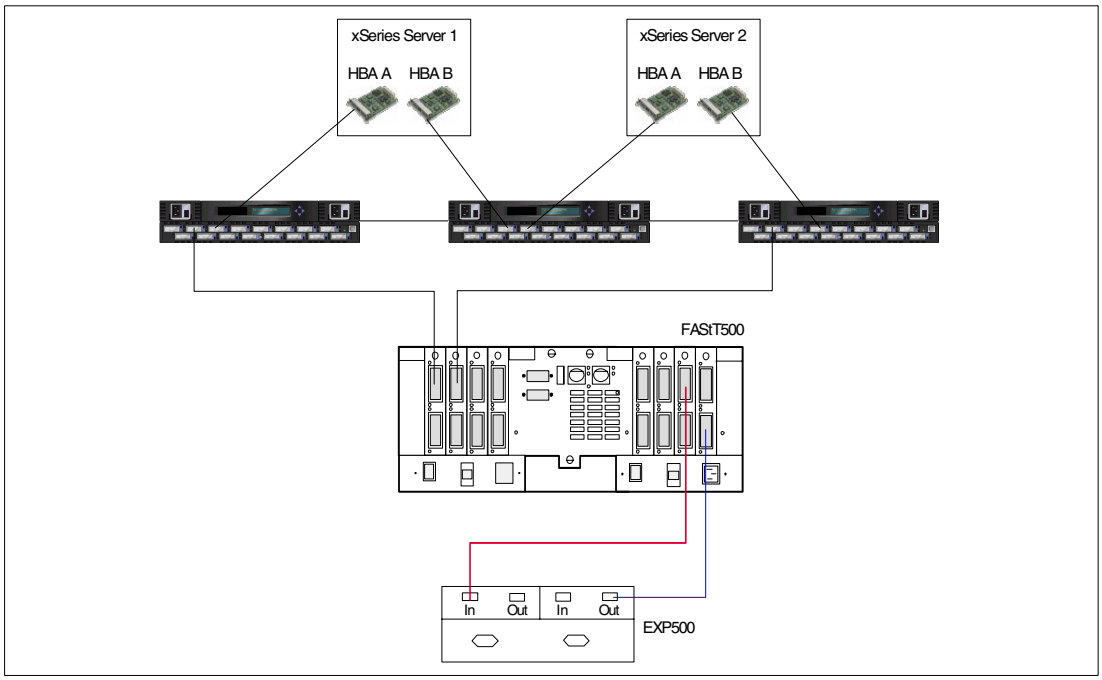

Figure 7-5 Cascaded switch fabric

8

Storage Expansion

No matter how much storage you have available, it seems there is always the demand for more. In this chapter we show how to expand the storage capacity of your FAStT500 Storage Server by adding EXP500 storage enclosures. In comparison with SCSI-based storage subsystems, Fibre Channel allows you to grow your storage capacity to levels that could not be achieved. To make sure that performance does not suffer as capacity increases, we also discuss increasing the drive-side bandwidth of a storage server.

8.1 EXP500 HotScale Technology

HotScale Technology allows us to scale a storage server while it is still in production serving disk resources to servers, thus increasing storage capacity with no downtime. By adding EXP500 storage enclosures to your configuration, you are increasing the total number of disk drive bays available to the storage server. It is then a simple task to add new disk drives and increase the total storage capacity of the FAStT500 Storage Server.

We now describe the steps we used to expand our existing EXP500 configuration. If you require information on other configuration options, please refer to the *FAStT500 Storage Server Installation Guide*, shipped with the product.

8.1.1 Steps to hot-add an EXP500

Our task is to expand our storage subsystem by adding an EXP500 storage enclosure while the storage subsystem is still performing disk I/O and serving disk resources to servers.

In our example, we have a single EXP500 attached to our FAStT500 Storage Server, and wish to expand our storage server by adding a second EXP500. The FAStT500 in our configuration has one pair of drive-side mini-hubs.

Important: Read this procedure carefully before attempting to make changes to your storage subsystem's EXP500 configuration. Failure to do so can mean interruption to I/O on your storage subsystem.

© Copyright IBM Corp. 2002

5. Disconnect the fiber cable from the drive-side mini-hub pair IN port on the FAStT500, at the OUT port on the right-hand ESM board of the existing EXP500. Do not disconnect any other cables from the EXP500. Reattach this cable to the OUT port on the right-hand ESM board of the new EXP500 (see Figure 8-3):

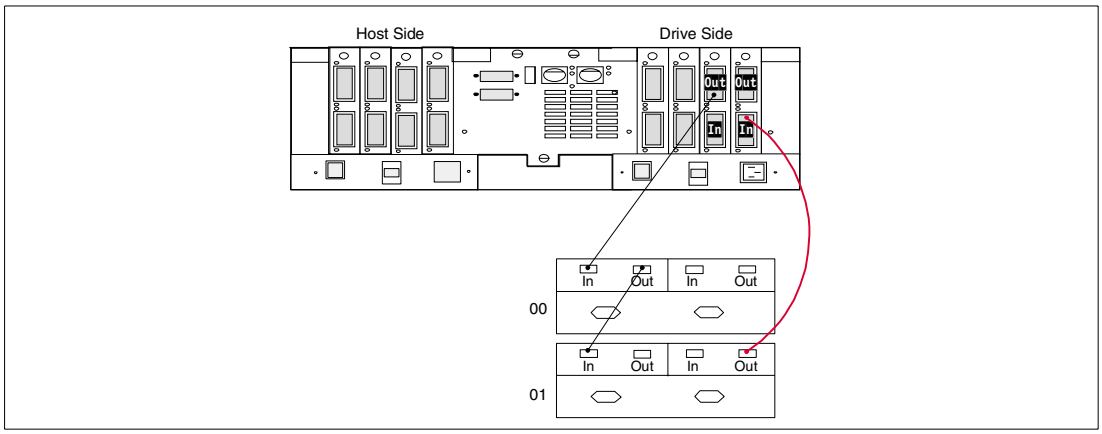

Figure 8-3 Connecting the new EXP500, step 2

6. Attach a new fiber cable from the OUT port on the right-hand ESM board of the current EXP500 to the IN port on the right-hand ESM board of the new EXP500.

   The FAStT500 Storage Server has now been expanded to two EXP500 storage enclosures, with no interruption to service (see Figure 8-4):

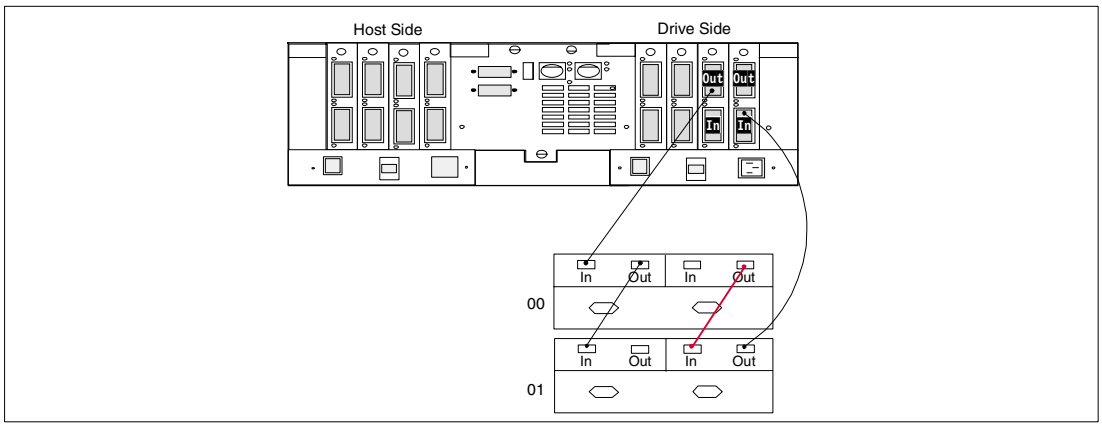

Figure 8-4 Hot Add EXP500 diagram 4

7. New disk drives can now be inserted into the second EXP500 and configured using Storage Manager.

This procedure can be simply extrapolated to adding multiple EXP500s simultaneously. Remember to disconnect only one live loop as we did in step 5. This is crucial, since, should you disconnect the other loop at the same time, your storage subsystem will no longer be able to perform I/O and you could lose data.

Using the steps as described, you are extending the live loop and breaking the redundant loop. You then reconnect the redundant loop after adding your new EXP500s.

8.2 Installing additional drive-side mini-hubs

Using two drive-side mini-hub pairs increases the drive-side bandwidth and you can balance the attachment of EXP500 enclosures to the FAStT500 Storage Server.

By installing another mini-hub pair into the FAStT500 Storage Server drive side we are able to expand the maximum number of attached EXP500 storage enclosures from 11 to 22, and increase the drive-side bandwidth of our storage server. Dividing the attachment of EXP500 storage enclosure between the two mini-hub pairs allows you to balance the load. For example, if you have 10 EXP500 storage enclosures, connecting five of them to one mini-hub pair and the other five to the other mini-hub pair splits the load. As you expand further, keeping the numbers of EXP500 units attached to each pair of mini-hubs equal (or close to equal) helps to ensure the bandwidth is used optimally. Of course, the workload placed on particular disks has an impact on the overall balance and also needs to be taken into account.

For more information please refer to the *FAStT500 Storage Server Installation Guide*.

8.3 Real examples

We close this chapter by reiterating the outstanding scalability of FAStT500 Storage Servers.

During the writing of this paper, we created a small storage server with a mesh fabric, which we used to illustrate the points we were making. Figure 8-5 shows the final configuration of this SAN:

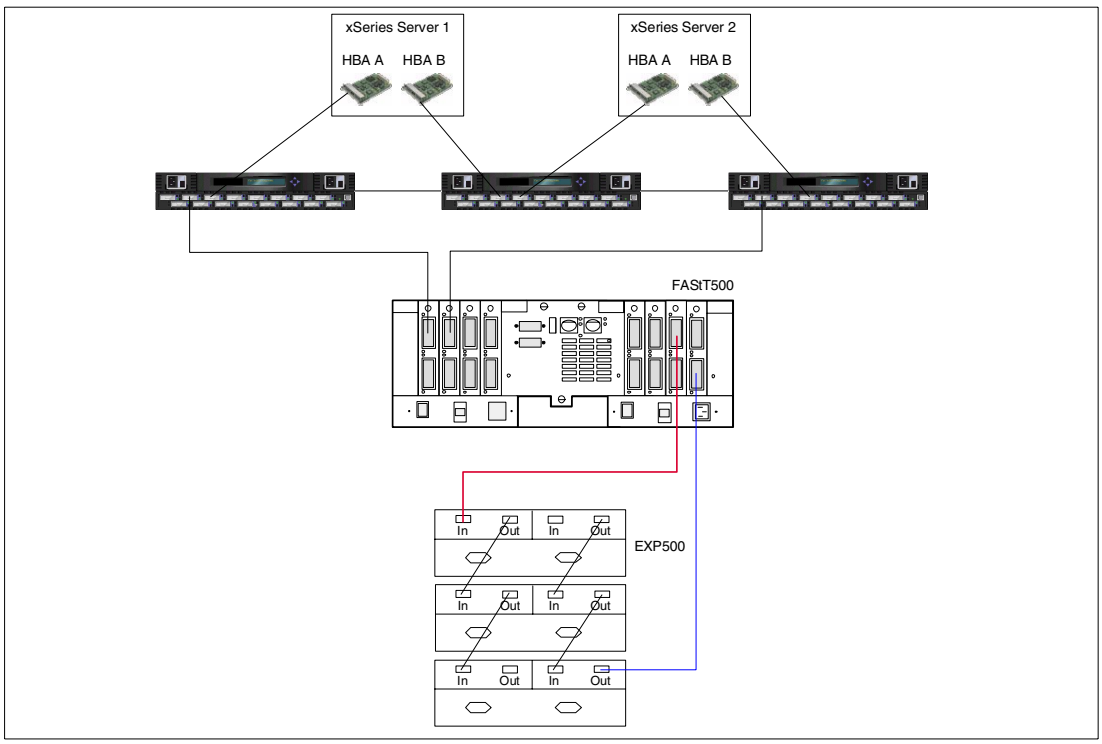

Figure 8-5 The SAN created while writing this Redpaper

A much bigger FAStT500 storage server is in live production use at Capespan (Pty) Ltd, South Africa. This subsystem was implemented by Andrew Le Crerar, one of the authors of this paper. The configuration currently uses nine EXP500 storage enclosures, with a total of 90 disk drives, for a total physical storage capacity of 3.2 TB. Five of the EXP500 storage enclosures are connected to one mini-hub pair and the remaining four enclosures are connected to the other mini-hub pair. The capacity can be further expanded with up to 13 additional EXP500 units and up to 130 additional disk drives.

Figure 8-6 shows the cabling of the nine EXP500 storage enclosures:

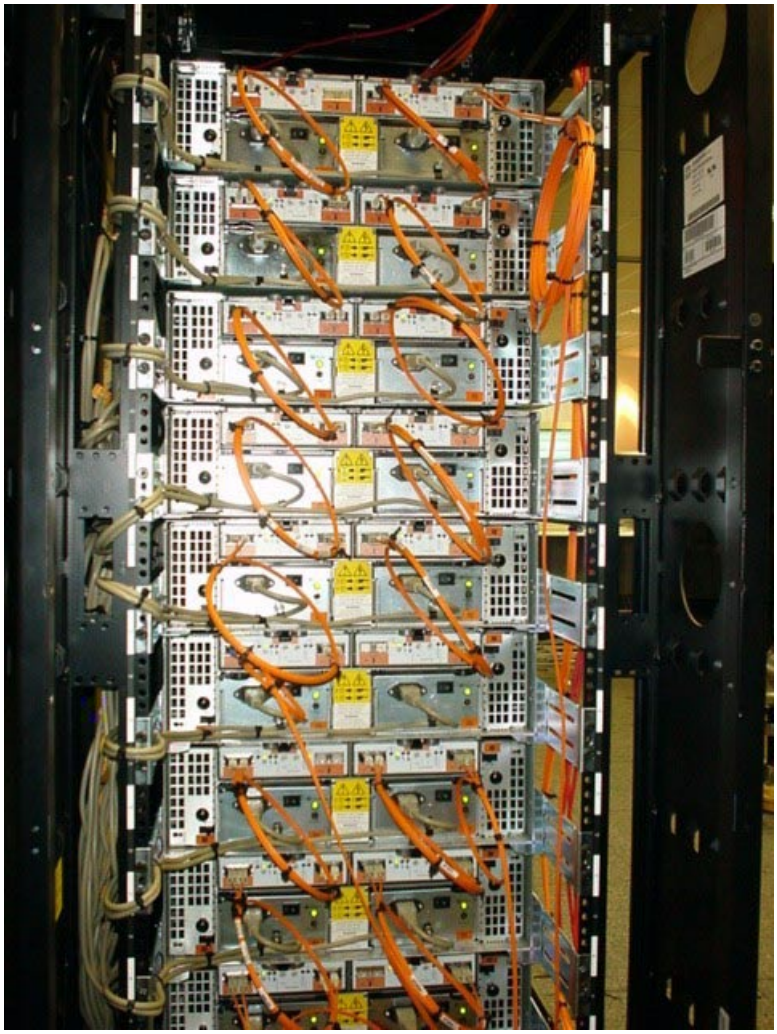

Figure 8-6 Photograph of FAStT500 cabling at Capespan (Pty) Ltd, South Africa

A

Creating heterogeneous SANs

One of the promised benefits offered by SANs is support for heterogeneous environments. In this appendix we demonstrate how this can be achieved by adding an xSeries server running Red Hat Linux to our existing Windows 2000-based SAN. You can use this as an example of the steps you might take to create your own heterogeneous SAN.

A.1 Pre-installation

The server we wished to add to our SAN is an xSeries 330 Server with the local disks attached to the on-board Adaptec SCSI controller and two FAStT HBAs installed. Red Hat Linux 7.1 was installed on this system. For our management client, we were testing the new Storage Manager Version 8, installed on one of our Windows 2000 machines. Similar steps apply if you wish to use Version 7.10.

The following software is required:

- Red Hat 7.1 installed, with the kernel at level 2.4.3-12. Details for upgrading the Red Hat Linux kernal are available at:

  http://www.redhat.com

- IBM FAStT Host Adapter Device Driver for Red Hat Linux Version 5.32.
- IBM FAStT Management Suite Java (MSJ) Diagnostic and Configuration Utility for Red Hat Linux Version 2.0.

A.2 Installation of FAStT HBA

The following will help you install and configure the FAStT Host Based Adapter. To physically install the adapters in your xSeries server, please see the installation guide supplied with the FAStT HBA.

© Copyright IBM Corp. 2002

A.2.1 FAStT HBA BIOS settings

When we installed the FAStT HBA in a Windows environment (see 3.2, “Host Bus Adapter device driver installation” on page 21), we set the adapter to the IBM defaults. For use with Linux, we have to modify some of these settings. It was found that changing these parameters improved Linux failover. The following steps will help you to configure the IBM FAStT Host Adapter.

1. Boot the xSeries Server with BIOS Version 1.76 or later diskette in the floppy drive. From the command prompt, run the following commands to update the adapter firmware (these are exactly the same as for the windows machine):

   ```
   flasutil /f /l
   flasutil /u.
   ```

2. Reboot the server.

Now that all the adapters in the system are at the latest BIOS levels we have to change some of the IBM default settings as stated previously.

3. During POST, press <ALT><Q> when you see the adapter’s BIOS screen, seen in Figure 3-1 on page 20.
4. In the Host Adapter Settings Menu, set the following values:
   - Loop Reset Delay = 8
   - Adapter Hard Loop ID = Enabled
   - Select a Hard Loop ID - 125 for example
5. In the Advanced Adapter Settings window, set the following values:
   - LUNs per Target = 0
   - Port down retry count = 12
   - Enable Target Reset = Yes

A.3 Support for the FAStT HBA in Linux

As the FAStT HBAs were installed during the installation of Red Hat Linux 7.1, the default kernel for FAStT Host Adapter was installed. This version does not work correctly with IBM FAStT_MSJ. You can either build the new driver from source or install the device driver.

A.3.1 Building the FAStT kernel from source

When Red Hat 7.1 installs, it has a downlevel kernel for the FAStT HBA. You must upgrade the driver and edit some files to make the adapter work correctly.

The kernel-headers and kernel-sources RPM packages need to be installed. These are available from the Red Hat support Web site.

1. To install the packages, issue this command:

   rpm -ivh kernel-headers*.rpm and **rpm -ivh kernel-sources*.rpm**

   Note: If a conflict is reported during the installation of these packages, the **--force** parameter has to be used.

Now we build the driver from source:

2. Obtain IBM FAStT Host Adapter Device Driver Version 5.32 for Red Hat Linux 7.1.

3. Create a directory in which to put the file:

   mkdir qla2x00

4. Change to the new directory and copy the device driver to it:

   cd qla2x00
   mcopy a:*.tgz

5. Now extract the source files from the archive:

   tar -xvzf *.tgz

6. Build the driver for the qla2200 from the source with the following command:

   make all SMP=1

   This makes both the qla2200.o and qla2300.o objects for the 2.4.3-12smp kernel.

7. The object files for the adapter now need to be copied to the /lib/modules/2.4.3-12smp/kernel/scsi directory. To do this, enter:

   cp qla2200.o /lib/modules/2.4.3-12smp/kernel/scsi

   Note: If you are prompted that you are going to overwrite qla2200.o respond Y or Yes.

A.3.2 Installing the qla2x00 device driver using RPM

Now we have to install the device driver using the following steps:

1. Run the **lsmod** command to see if any drivers for the qla2x00 are currently loaded. Your output should be similar to Example 8-1.

Example 8-1 lsmod command output

```
[root@th71 /root]# lsmod
ModuleSizeUsed by
qla2x002255525
autofs118081(autoclean)
eepro100172321(autoclean)
ipchains416321(unused)
usb-ohci524161(unused)
usbcore524161(usb-ohci)
aic7xxx1363363
sd_mod117447
scsi_mod986243[qla2x00 aic7xxx sd_mod[
[root@rh71 /root]#
```

2. If the IBM FAStT Host Adapter was physically installed during the installation of RedHat 7.1 then the default driver will need to be unloaded using the following command:

   modprobe -r qla2x00

3. Install the IBM FAStT Host Adapter Device Driver Version 5.32 for RedHat Linux 7.1 with the command:

   rpm -ivh qla2x00-5.32-1-i686.rpm

4. If an older version of the driver is present, you need to add the **--force** parameter to the command:

   rpm -ivh --force qla2x00-5.32-1-i686.rpm

A.3.3 Edit modules.conf

The next step is to edit modules.conf so that the driver can see more than one LUN, with options scsi_mod max_scsi_luns=32.

1. Open the module.conf file for editing. We used the vi editor:

   vi /etc/modules.conf

2. Press Insert, move the cursor to end of the last line, press Enter, and add the following line:

   ```
   options scsi_mod max_scsi_luns=32
   ```

   The file should now look something like Example 8-2:

Example 8-2 modules.conf file after editing

```
alias eth1 eepro100
alias scsi_hostadapter aic7xxx
alias scsi_hostadapter qla2200
alias usb-controller usb-ohci
options scsi_mod max_scsi_luns=32
```

3. Press Esc, and type `wq` on the command line to write the changes to the file and quit.
4. Press Enter.
5. As a result of updating /etc/modules.conf, it is necessary to update any dependencies/lib/modules/2.4.7-10/modules.dep by running the command:

   depmod -a

6. Reboot the server.
7. Check if the HBA kernel module is loaded with the command:

   lsmod qla2200

8. If qla2x00 is not loaded, install it with the command:

   modprobe qla2x00

   Note: **modprobe** must be used to load the options string. **insmod** will load the driver without reading the options string and should not be used.

A.4 Create a new LUN and assign to Hosts

Using Storage Manager, which we have already installed on the Windows 2000 management machine, we can create and assign LUNs to the new Linux server.

We took the following steps to achieve this. Note that these steps are very similar to those already seen in 5.3, “Using the Storage Manager Client” on page 54, so we omit some of the fine detail:

1. Storage Manager starts at the Enterprise Management window, which displays the storage servers that you are able to manage, as shown in Figure A-1:

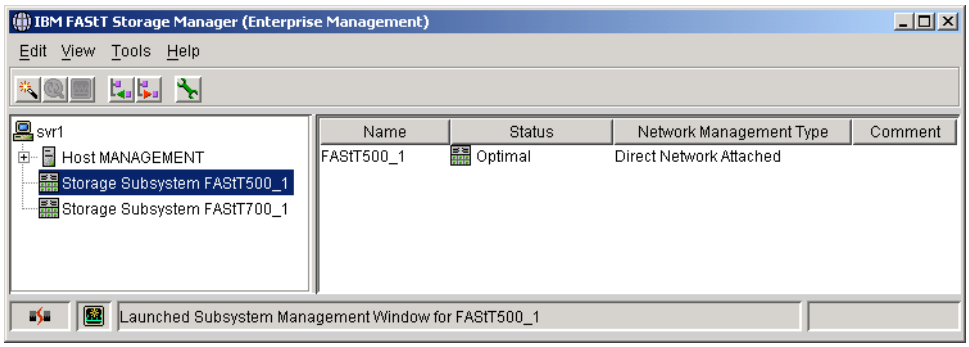

Figure A-1 Selection of FAStT RAID Controller

2. Right-click **Storage Subsystem FAStT500_1** and select **Manage Device** to display the Subsystem Management window, which shows you the current configuration and free capacity, as shown in Figure A-2.

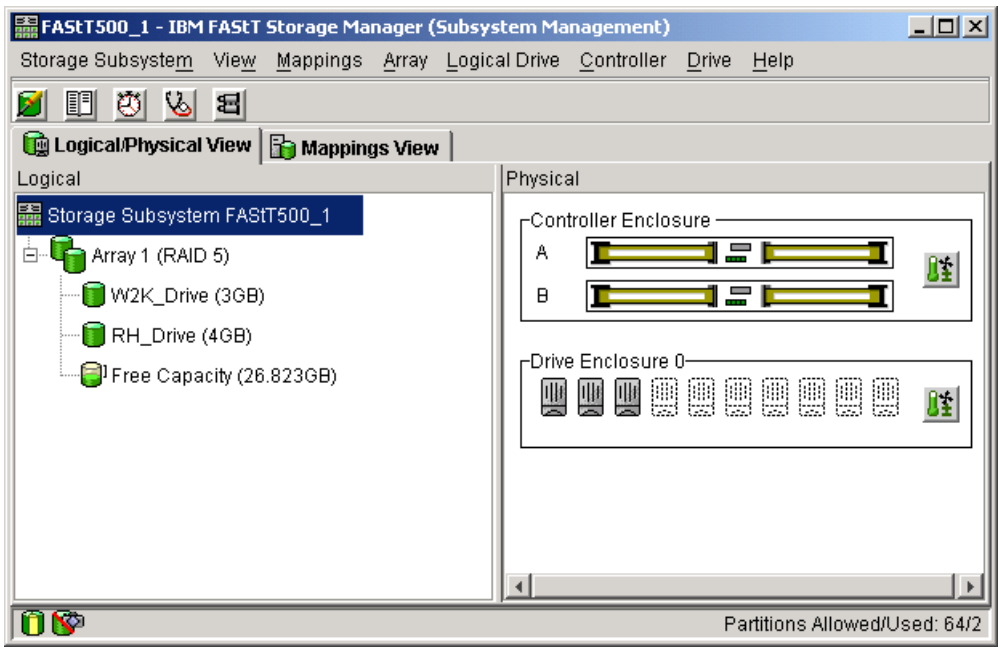

Figure A-2 Subsystem Management window

3. Right-click **Free Capacity** and select **Create Logical Drive...** to create a logical drive for the new xSeries Linux server, as shown in Figure A-3.

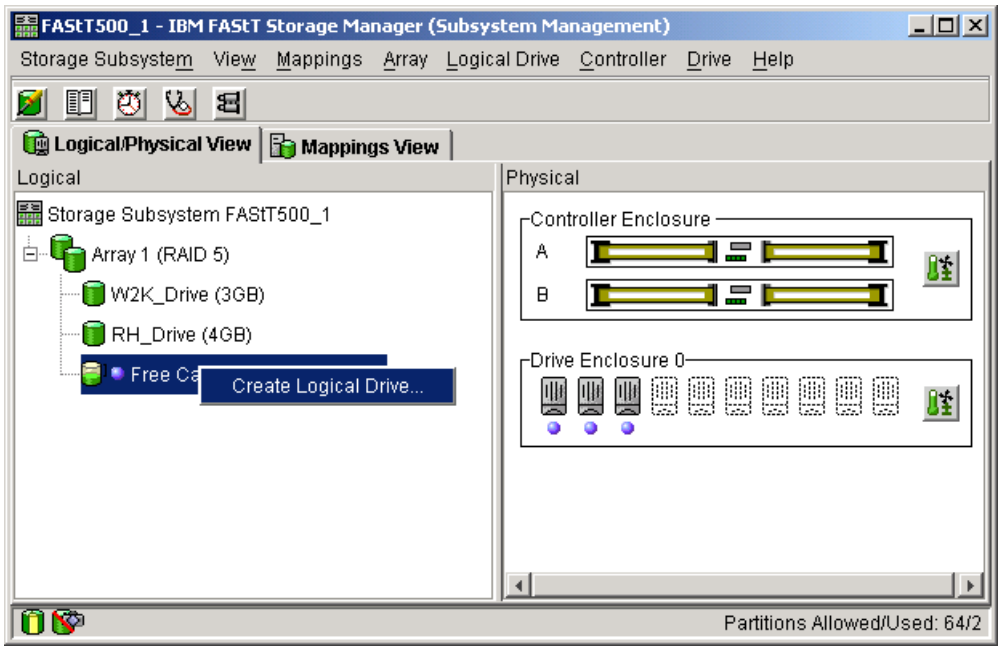

Figure A-3 Create logical drive

4. Since our server (host) is running Linux, we must define a host type. We set this to be Linux to match our host as shown in Figure A-4:

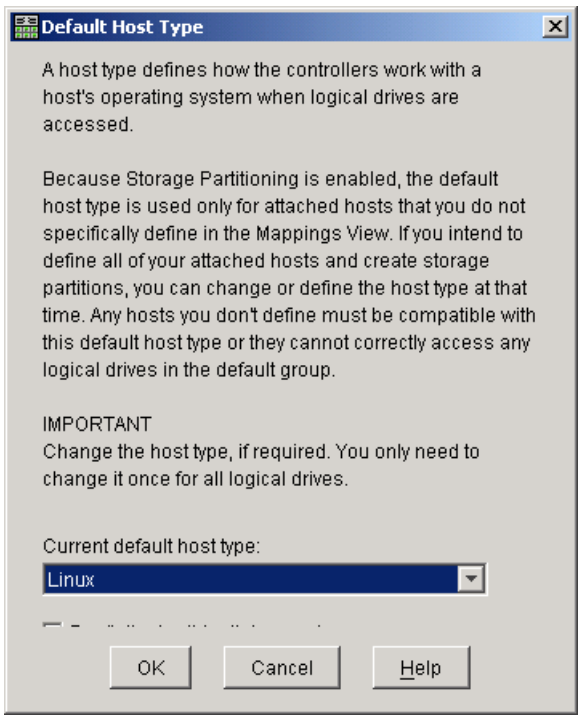

Figure A-4 Default Host Type window

5. Click **OK** and you are prompted for a password before you are allowed to proceed, as shown in Figure A-5:

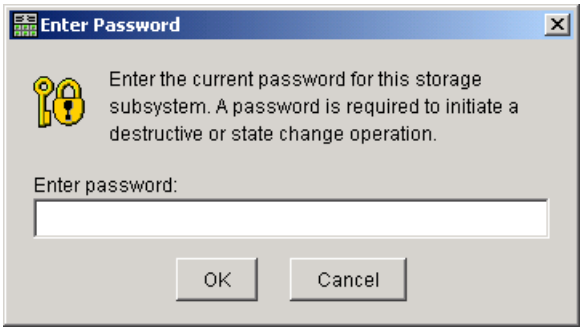

Figure A-5 Password prompt

6. After entering the correct password, you can proceed with creating a logical drive, as shown in Figure A-6:

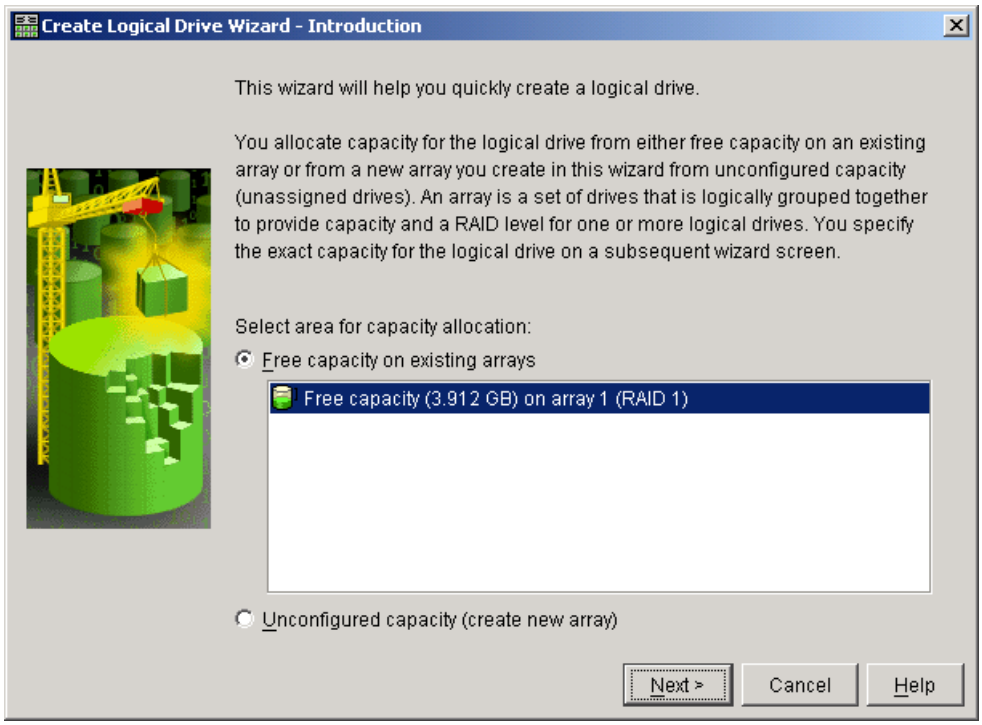

Figure A-6 Create logical drive

7. For our new logical drive, we used some of the free drive capacity in Array 1. Once you have selected your capacity type, click **Next**.
8. Now we have to specify the logical drive parameters. This includes the drive size and a volume name. The name should be clear and follow a good naming convention. We allocated 2 GB and called our volume `SVR2_Volume`, as shown in Figure A-7:

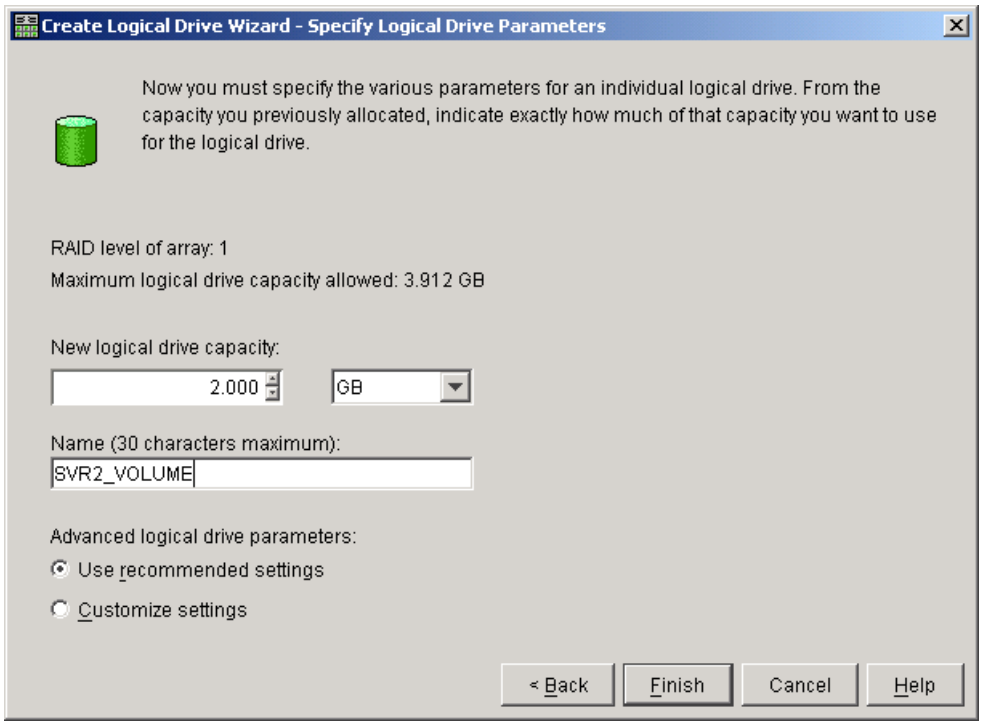

Figure A-7 Specify logical drive parameters

9. Now that we have defined a logical drive it is time to map this to the Host. Click the **Mappings View** tab, which can be seen in Figure A-2 on page 97, then right-click **Default Group** and create a new Group, which we called `SVR2`.
10. Right-click **SVR2** and select **Define New Host Port...**. Use the WWN information for the adapters and give a name that is meaningful. For example, in our configuration, the two FAStT HBAs were named `SVR2_HBA_A` and `SVR2_HBA_B` respectively.
11. Open **Undefined Mappings** and right-click on the new volume (**SVR2_Volume**).
12. Map this logical drive to the Host SVR2 group. Once this is complete, the new logical drive is accessible by Host SVR2, as shown in Figure A-8:

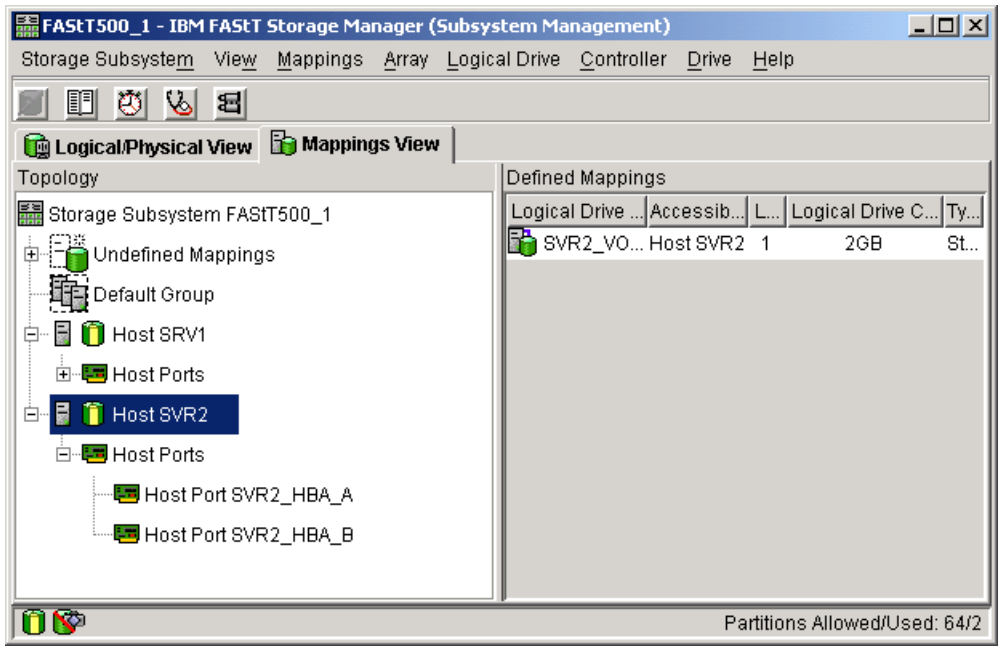

Figure A-8 Configured Linux server with 2 GB drive

A.5 Zoning

We now have created some disk space for the Linux xSeries server. The next step is to update the zone configuration for the SAN to allow the Linux xSeries servers adapters to see the FAStT500 RAID Controllers. In our case we have to add the new adapters to the zones already created in 4.6, “Zone configuration” on page 37. To do this we performed the following steps:

1. Open a Web browser and enter the IP address of the switch as the URL.
2. Enter the Zone Administration window.
3. Create aliases for the two adapters in this server. Our server name is SVR2 and the adapters are called SVR2_HBA_1 and SVR2_HBA_2. This is shown in Figure A-9:

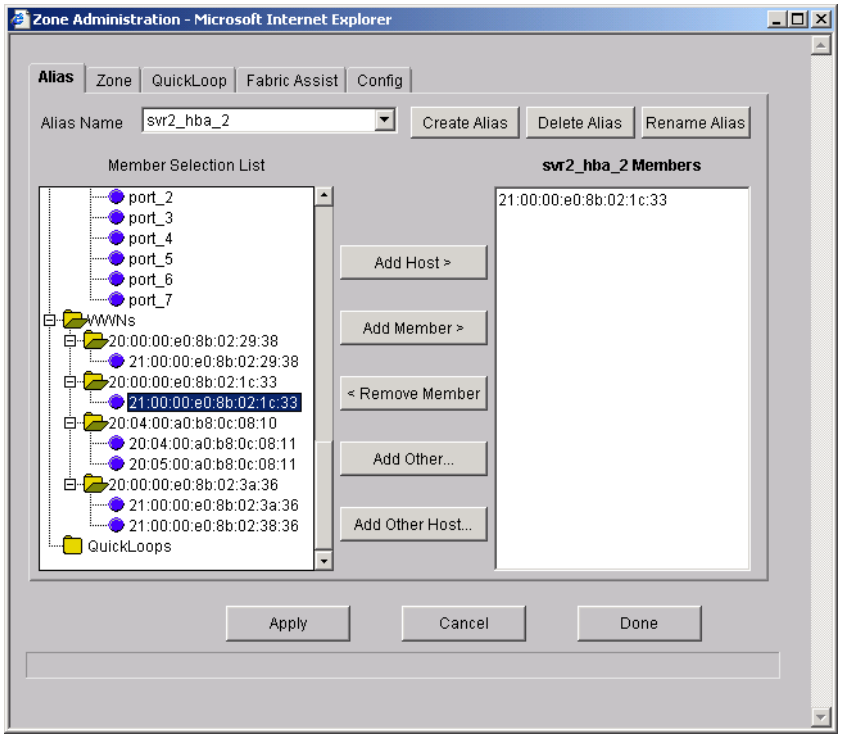

Figure A-9 Create Alias

4. Having created aliases, we now add these to our existing zones for Contoller A and Controller B, as shown in Figure A-10:

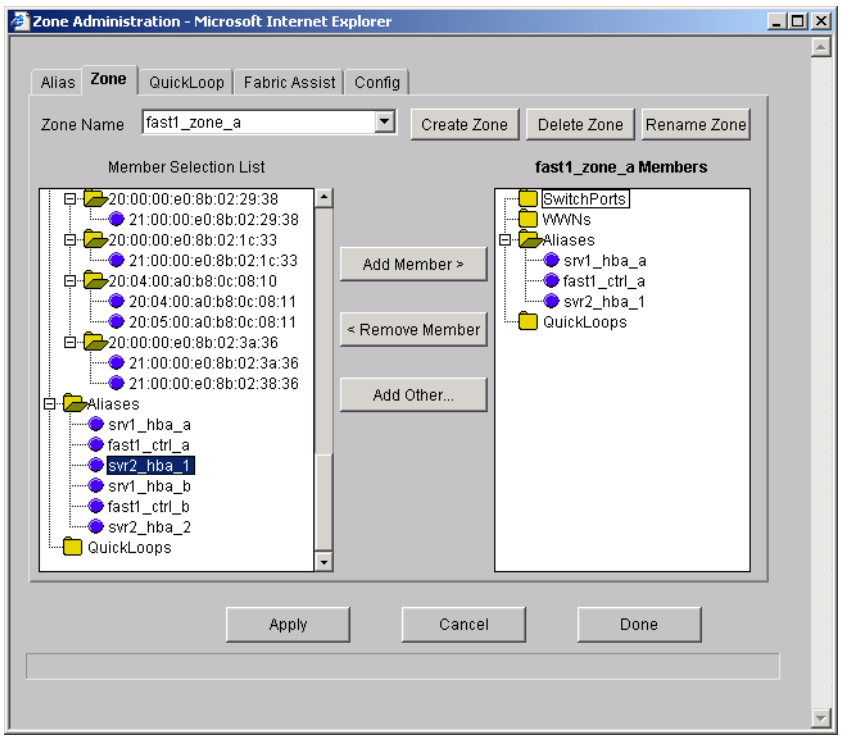

Figure A-10 Zone configuration

5. After updating the zones, it is a good practice to ensure that they are enabled. Select the **Config** tab and check **Enable**, as shown in Figure A-11:

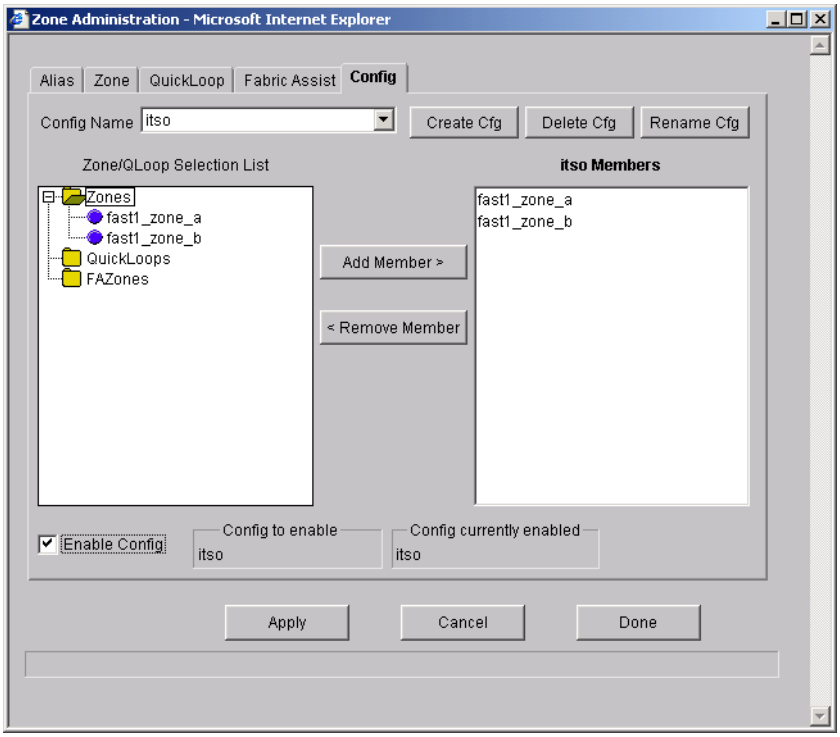

Figure A-11 Zone configuration enabled

6. Click **Apply** and **Done**.

Now we have completed the zoning administration required to add the additional server.

A.6 Installing FAStT_MSJ

The Management Suite Java is used to configure Linux Host Bus Adapter failover.

1. First verify the driver version that is currently loaded with this command:

   cat /proc/scsi/qla2200/1

   The output should look similar to Example 8-3:

Example 8-3

```
QLogic PCI to Fibre Channel Host Adapter for ISP2100/ISP2200/ISP2200A:
Firmware version:2.01.34
Driver version 5.32
Entry address = d083a060
HBA: QLA2200 , Serial# A90923
```

2. In a command shell within X-Windows, change to the directory in which you extracted the FAStT_MSJ package and run this command:

   chmod 700 FAStTMSJ_install.bin

3. To start the installation process, enter:

   ./FAStTMSJ_install.bin

4. The Install window shown in Figure A-12 appears:

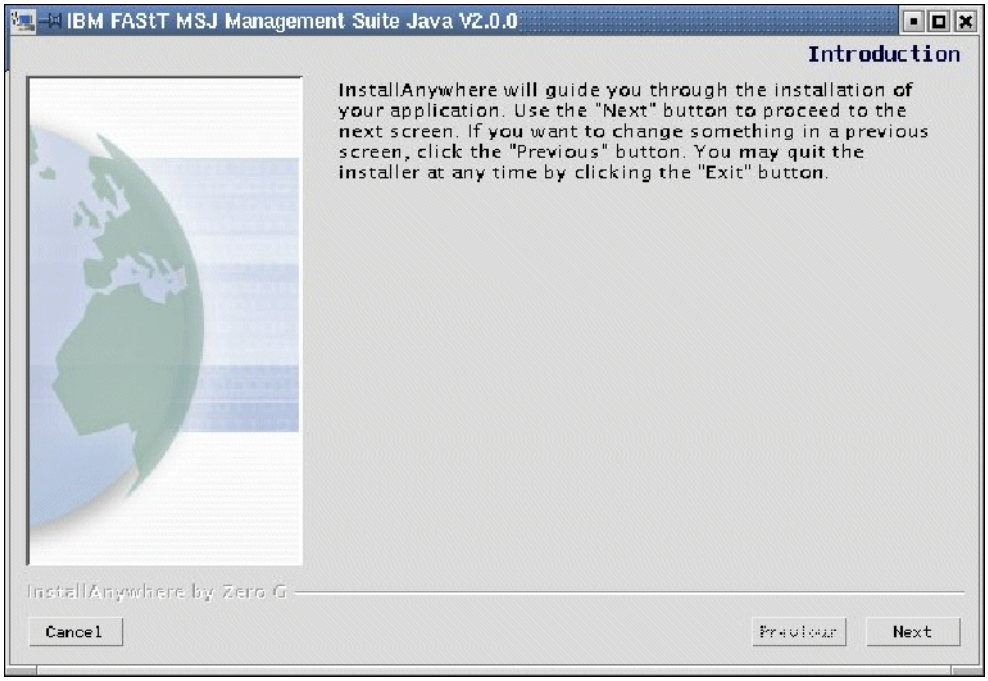

Figure A-12 Install window

5. Click **Next** to display the Choose Product Features window, as shown in Example A-13:

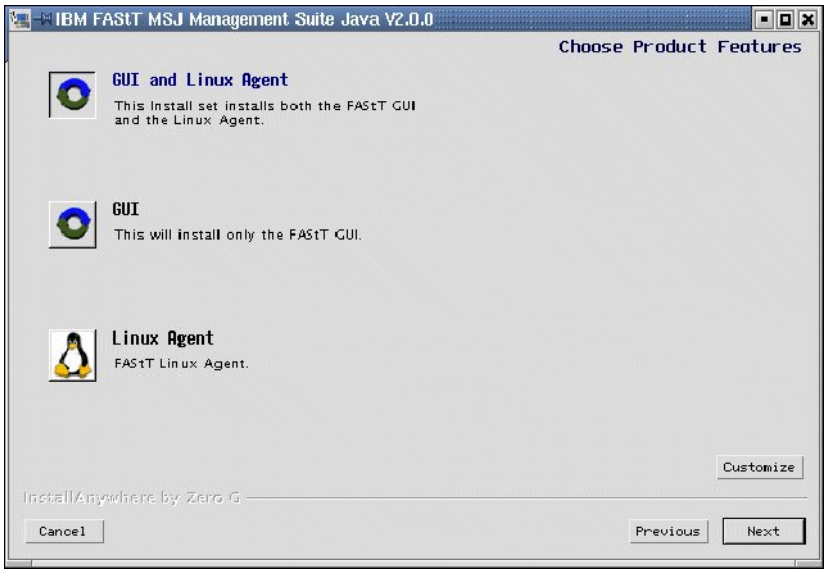

Figure A-13 GUI and agent selection

6. Select **GUI and Linux Agent** and click **Next** to display the Choose Install Folder window.
7. Accept the default and click **Next**.
8. The install process will complete without further interaction. Click **Done** to close the Install Complete window.

IBM FAStT_MSJ has been successfully installed.

A.7 Configuring FAStT_MSJ

We now need to configure the FAStT_MSJ for failover.

1. Start the qlremote agent. This agent needs to run in the background. At a command prompt run **qlremote**.

   If the driver is not the correct driver for the FAStT_MSJ package you will get the termination error shown in Example 8-4.

Example 8-4 qlremote error

```
[root@rh71-2 Redhat71]# qlremote
Debug: OSS initialized...
Debug: Core initialized...
Debug: RPC initialized...
Debug: CorePollingLoop() starting...
Error: SDOpenDevice (0, xx) failed (0x9) (System Error 9.)
Error: **NO** compatible HBAs found during scan - Terminating...
Debug: Signal (15) caught...
Debug: Shutting Down...
Debug: Deallocating Core device list...
Debug: Deallocating OSS resources...
Debug: Deallocating RPC resources...
Terminated
[root@rh71-2 Redhat71]# Debug: CorePollingLoop() ending...
```

   To verify the driver version that is currently loaded, enter the command:

   cat /proc/scsi/qla2200/1

2. With the qlremote agent loaded, open another command window to start IBM FAStT_MSJ with the command:

   /opt/IBM_FAStT_MSJ/FAStT

   This will open a GUI showing you the HBA view, as shown in Figure A-14:

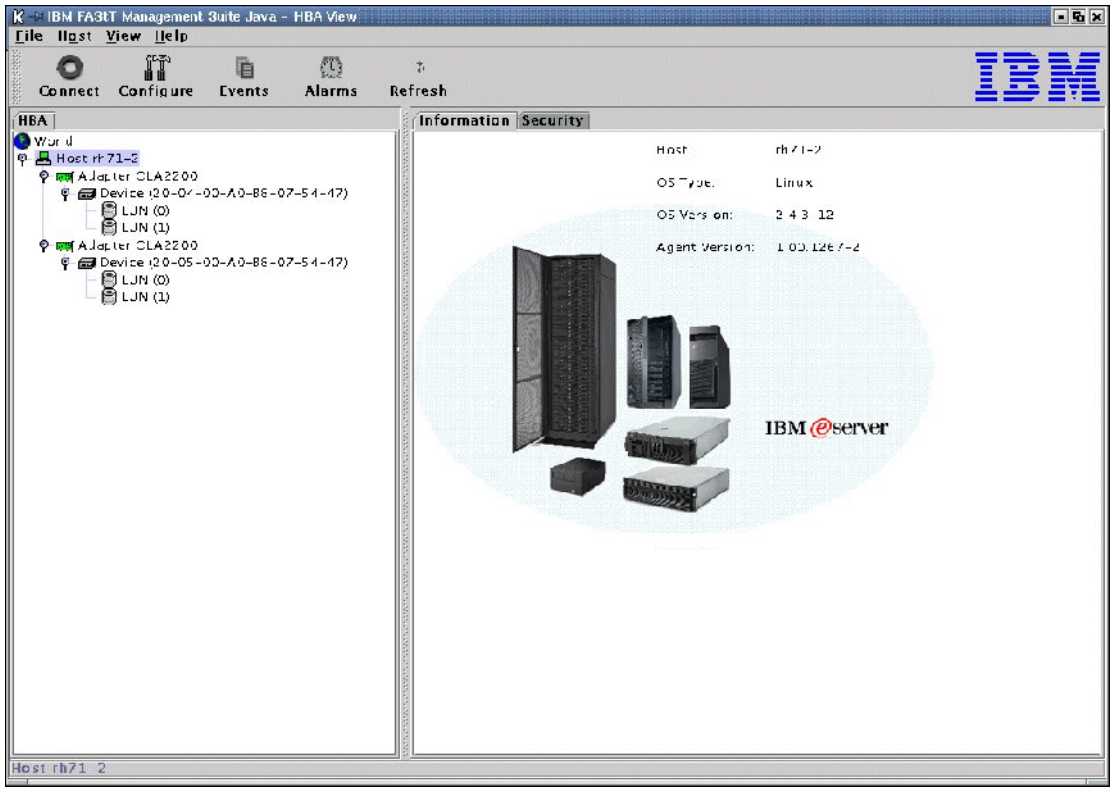

Figure A-14 FAStT_MSJ GUI - HBA view

3. Click **Connect**, enter the IP address of the local server and press Enter.
4. Click **Configure** then select **Device->Autoconfigure** from the menu.
5. Select the LUNs you wish to configure. In our case, we selected LUN 0 and LUN 1.
6. Select **Load Balance**.
7. Select LUNs and click **OK**.
8. After the configuration is complete, click **OK**. This will take you back to the Fibre Channel configuration window. Click **Apply**.
9. You will be prompted for a password. The password is **config**.
10. Click **OK** to accept the refresh.

    The LUN configuration window shown in Figure A-15 is displayed:

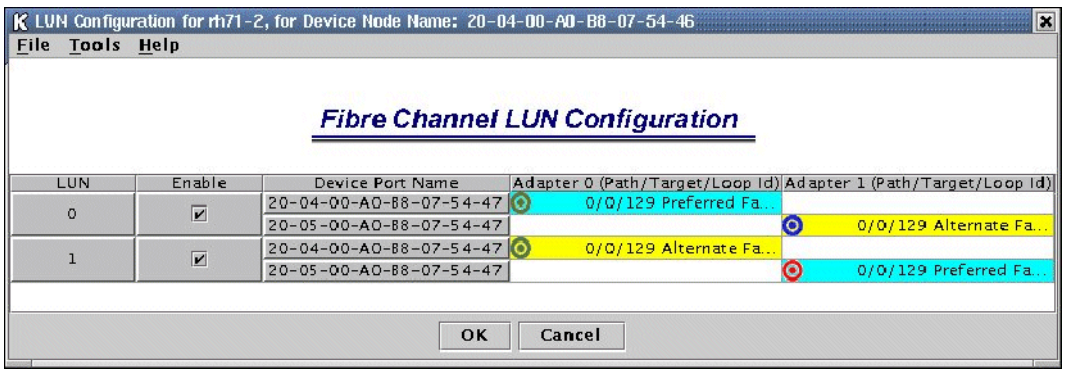

Figure A-15 LUN configuration window

Important: You must now unload and reload the driver using **modprobe**.

11. Unload the driver with the command:

    ```
    modprobe -r qla2200
    ```

12. Reload the driver with the command:

    ```
    modprobe qla2200
    ```

13. Now you must rebuild the boot image so that the correct driver parameters are loaded when the server restarts. Enter the command:

    ```
    mkinitrd -f /boot/newinitrd-image 2.4.3-12
    ```

14. Reboot and verify that the correct driver and multi-path configuration is loaded. After rebooting the server, execute this command from a prompt:

    ```
    cat /proc/scsi/qla2200/1
    ```

A.8 Activating the Linux devices

Examine /var/log/messages and check that the new drives are identified.

Also if you go into System and look at Hardware Properties you should see a new hard drive. In our case we saw the new 2 GB drive we created.

1. Use **fdisk** to create a partition on the new drive.
2. Run **mkfs** to create a file system in the new partition. For example:

   ```
   mkfs /dev/sdc1
   ```

3. Run **mkdir** to create a directory. For example:

   ```
   mkdir /sandisk
   ```

4. Run **mount** to mount the new drive. For example:

   ```
   mount /dev/sdc1 /sandisk
   ```

5. Edit fstab to update the filesystem information. For example, add the following line, using the vi editor:

   ```
   /dev/sdc1 /sandiskdefault1 2
   ```

B

FAStT Management Suite Java, FAStT Storage Manager 8 and SANavigator

In this appendix we provide more information about the installation and use of the FAStT Management Suite Java. In addition, although FAStT Storage Manager 8 and SANavagator had not been released, we wanted to give you an overview of some of the new features and functionality that should be available.

B.1 FAStT Management Suite Java 2.0.0

This tool is used to manage the Host Bus Adapters in your xSeries SAN environment and is the replacement for FAStT Check. It comprises two components:

- An agent that is installed on all servers that have one or more Host Bus Adapters.
- The GUI that can be installed on any machine and which is used to manage the Host Bus Adapters installed in the servers.

The agent is supported on the following operating systems:

- Windows NT 4.0
- Windows 2000
- Red Hat Linux 7.1 upgraded to Kernel 2.4.3-12
- Novell NetWare 5.1

The GUI is supported on the following operating systems:

- Windows NT
- Windows 2000
- Red Hat Linux 7.1 upgraded to Kernel 2.4.3-12

© Copyright IBM Corp. 2002

The latest version of FAStT MSJ is available from:

http://www.pc.ibm.com/support

- On this Web page, from the Browse list on the right-hand side of the page, click **Servers**.
- On the next page, in the Family listbox, select **Fibre Channel Solutions**.
- On the left-hand navigation bar, click **Downloadable files**.
- In the Downloadable files by category listbox, select **Fibre Channel Solutions** to get a listing of all Fibre Channel files.
- Locate the file in the list and click its name to go to the download page. In our example, the file was called Fibre Channel Solutions - IBM FAStT Management Suite Java (MSJ) Diagnostic and Configuration Utility Version 2.0. Separate package files contain the installation executable for the supported operating systems.

B.1.1 Installing the FAStT Management Suite Java software

In this section we cover the steps to install FAStT MSJ on a Windows 2000 system.

Before installing FAStT MSJ, uninstall FAStT Check and the Intel DMI 2.0.

FAStT MSJ supports the following two configurations:

- Stand-Alone

  If you wish to manage the Host Bus Adapters in a stand-alone system then you would install both the GUI and Agent on the local system.

- Networked

  If you wish to manage the Host Bus Adapters on a remote system then you would install only the agent on the remote system and the GUI on a remote management system. A remote management system may contain Host Bus Adapters.

Follow these steps to install FAStT MSJ:

1. Unpack the downloaded package file to a path that you specify, by running the self-extracting executable. In our example the file name was 33p2505.exe.
2. Install the FAStT MSJ by running the installation file unpacked in step 1. In our example the file name was FAStTMSJ_Install.exe.
3. The InstallAnywhere status window is displayed while the application prepares to install, followed by the FAStT MSJ splash screen. When the Introduction window is displayed, click **Next** to continue.
4. The Choose Product Features window is displayed, as shown in Figure B-1:

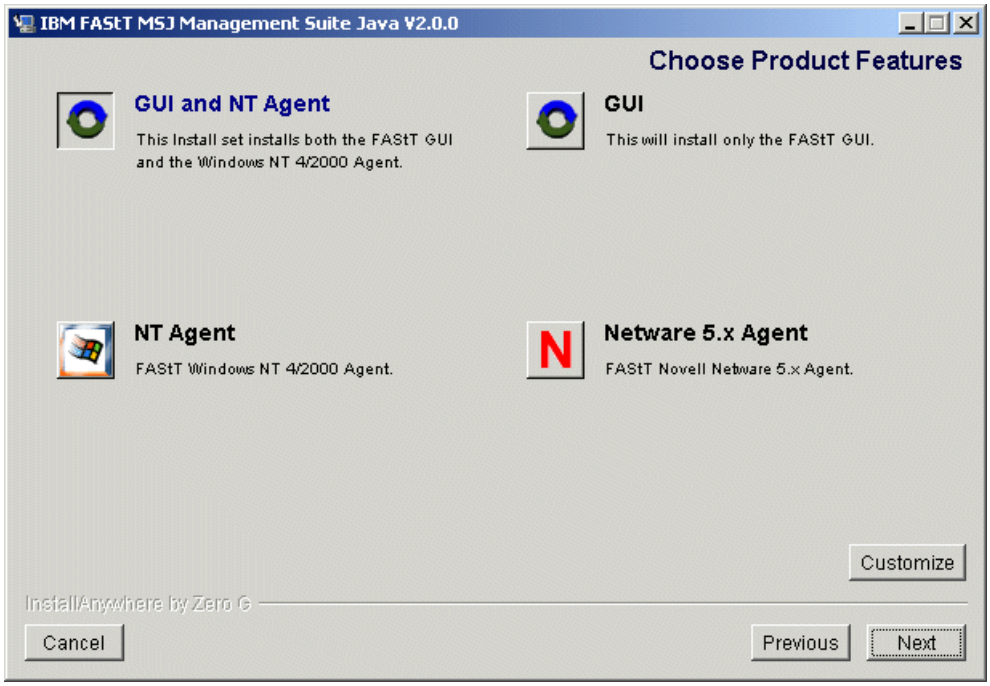

Figure B-1 IBM FAStT Management Suite Java install

5. Click the **GUI and NT Agent** icon and then click **Next**.

   We decided to install both the GUI and agent on the same machine in order to simplify this example. In reality, you would typically install just the GUI on a management workstation and the agent on machines containing Host Bus Adapters.

6. The Information window follows; click **Next**.
7. The Choose Install Folder is displayed. Specify your install path and click **Next**. In our example we left the path at the default.
8. The Select Shortcut Profile window is displayed. Select **All Users Profile** and click **Next**.
9. The Create Desktop Icon window is displayed. Check the **Create desktop icon** checkbox, and click **Install**.
10. The Installing window is displayed and shows the ongoing status of the installation process. When everything has been installed, the Install Complete window is displayed. Click **Done** to end the installation.

B.1.2 Using the FAStT Managment Suite Java GUI

The FAStT Management Suite Java GUI is used to manage servers running the FAStT MSJ agent. With this tool, you are able to view and make changes to the configuration of Host Bus Adapters in the server running the agent. You are also able to perform diagnostics tests and check the performance of the Host Bus Adapters. In this section we provide just a brief introduction to the use of this software. Please refer to the product documentation and online help for more information.

Follow these steps for a brief tour of the FAStT MSJ GUI:

1. Start the GUI by double-clicking the FAStT MSJ shortcut located on the desktop.
2. The IBM FAStT Management Suite Java - HBA View window is displayed (see Figure B-2):

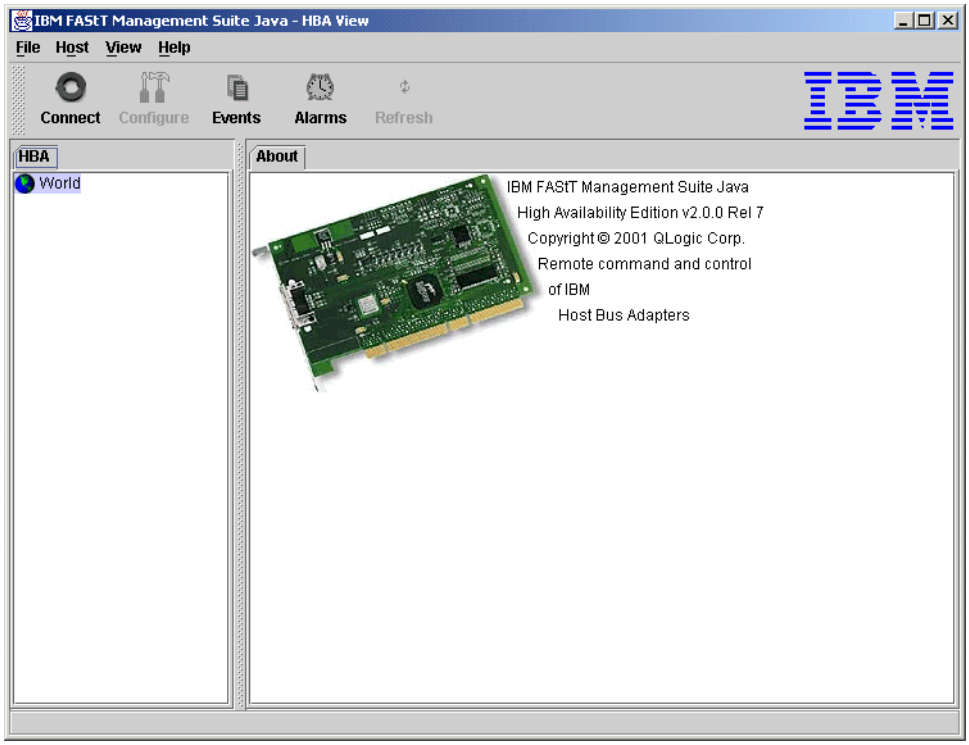

Figure B-2 IBM FAStT Management Suite Java - HBA View window

3. This is the main window in the GUI. To connect to a server agent click the **Connect** icon and the Connect to Host window is displayed (see Figure B-3):

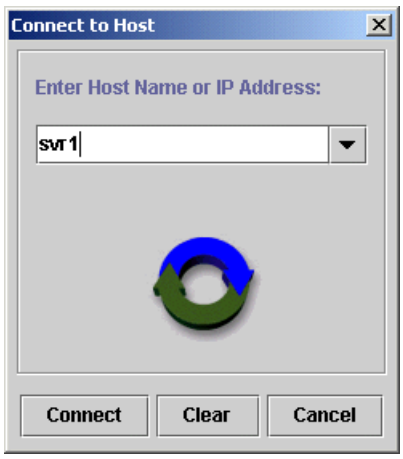

Figure B-3 FAStT MSJ - Connect to Host window

4. Enter the name (we used svr1) or the IP address of the server to which you wish to connect, and click **Connect**.
5. The server is added to the left-hand side pane of the main window, on the HBA tab. Host Bus Adapters installed in the server appear below it.
6. Click **Host svr1** to see two tabs displayed in the right-hand pane. They are:
   - Information, containing basic information about the currently connected server.
   - Security, containing security settings for the connected agent.

   The tabs are shown in Figure B-4:

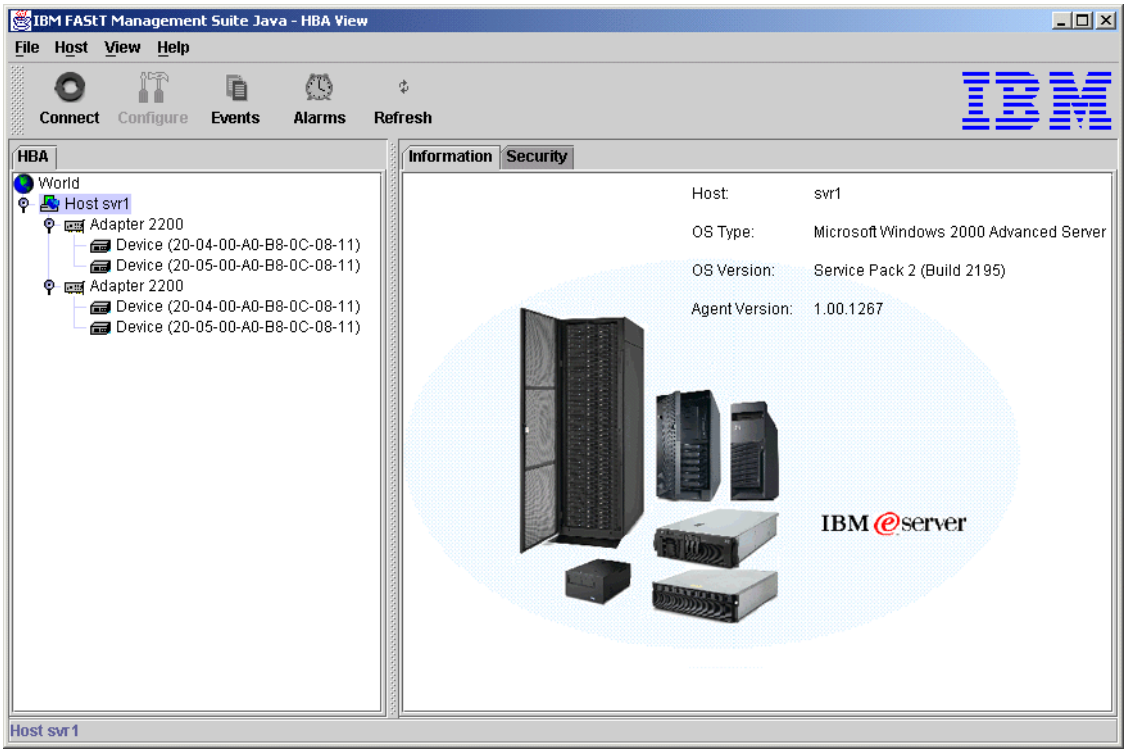

Figure B-4 Host selected within IBM FAStT Management Suite Java - HBA View window

7. Click one of the Host Bus Adapters in the left-hand pane and the right-hand pane changes. There are now seven tabs as shown in Figure B-5:

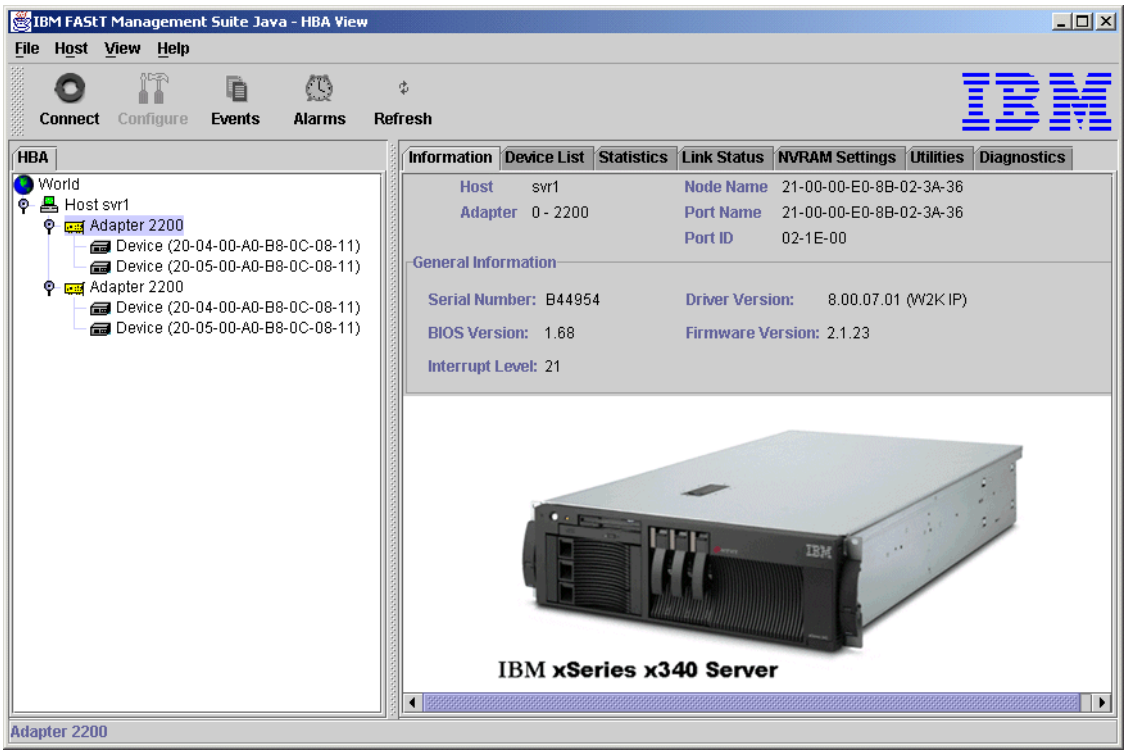

Figure B-5 HBA selected within IBM FAStT Management Suite Java - HBA View window

The following tabs are displayed:

- Information, which displays general information about the server and Host Bus Adapter.
- Device List, which displays the devices currently available to the Host Bus Adapter.
- Statistics, displaying a graph of the performance and errors on the Host Bus Adapters over a period.
- Link Status, displaying the current link status.
- NVSRAM Settings, which displays the current settings and allows you to make remote configuration changes to the NVSRAM.
- Utilities, which allow you to update the flash and NVSRAM remotely.
- Diagnostics, which allow you to run diagnostic tests remotely.

B.2 FAStT Storage Manager 8.0

This section discusses the new verison of Storage Manager that was not released at the time we wrote this paper. As the details are subject to change until formal release of the code, we provide an introduction to the main new features rather than a detailed exploration.

Important: Remember to read all documentation that comes with an upgrade before attempting the upgrade. This can be critical to the success of the upgrade.

We start by discussing FlashCopy and give step-by-step details for creating a FlashCopy logical drive, and then assign this logical drive to a host using the new storage partitioning view of Storage Manager 8. The other new feature we discuss is the ability to expand a logical drive and extend a partition using Windows 2000 dynamic disks.

The base functionality of the previous version of Storage Manager remains unchanged with the exception of a few minor GUI adjustments to buttons and views.

B.2.1 Creating a FlashCopy logical drive

With the ability to perform FlashCopy, you are able to take a point-in-time image of a logical drive. This can then be mounted onto any server in the SAN. The key thing to note is that a FlashCopy is not a physical copy of data from one logical drive to another. A FlashCopy operation tracks changes made to the original logical drive after the FlashCopy is created.

What this means is that, at the point at which you initiate the FlashCopy, a new logical drive is created, which appears and acts like an exact copy of the original. What really happens is that pointers to information on the original disk are mapped to the new FlashCopy drive. Now, when data is changed on the original logical disk, the old information is written to the FlashCopy while the original disk is updated with the new information.

A FlashCopy, therefore, takes only a few seconds even for large disks. All physical copy operations occur after the FlashCopy has been created, and only for data that is changed. This means that you always have the information as it was at that point-in-time. Obviously the longer you keep a FlashCopy logical drive, the more space it will use, as more and more of the data on the original disk is updated.

FlashCopy can be used to achieve many different tasks. A particularly useful example would be taking a FlashCopy of a logical disk containing a database, and then mapping it to another server to perform a backup, while the original logical drive and database continue to function.

Depending upon the application, a brief pause in operation may be required while the FlashCopy is performed. For example, it may be necessary to quiesce a database to ensure that the data held within it is consistent while the copy is made

The following steps create a FlashCopy with Storage Manager 8:

1. Within the Subsystem Management window, right-click the logical drive you wish to copy (**NewDrive** in our example) and select **Create FlashCopy Logical Drive** (see Figure B-6):

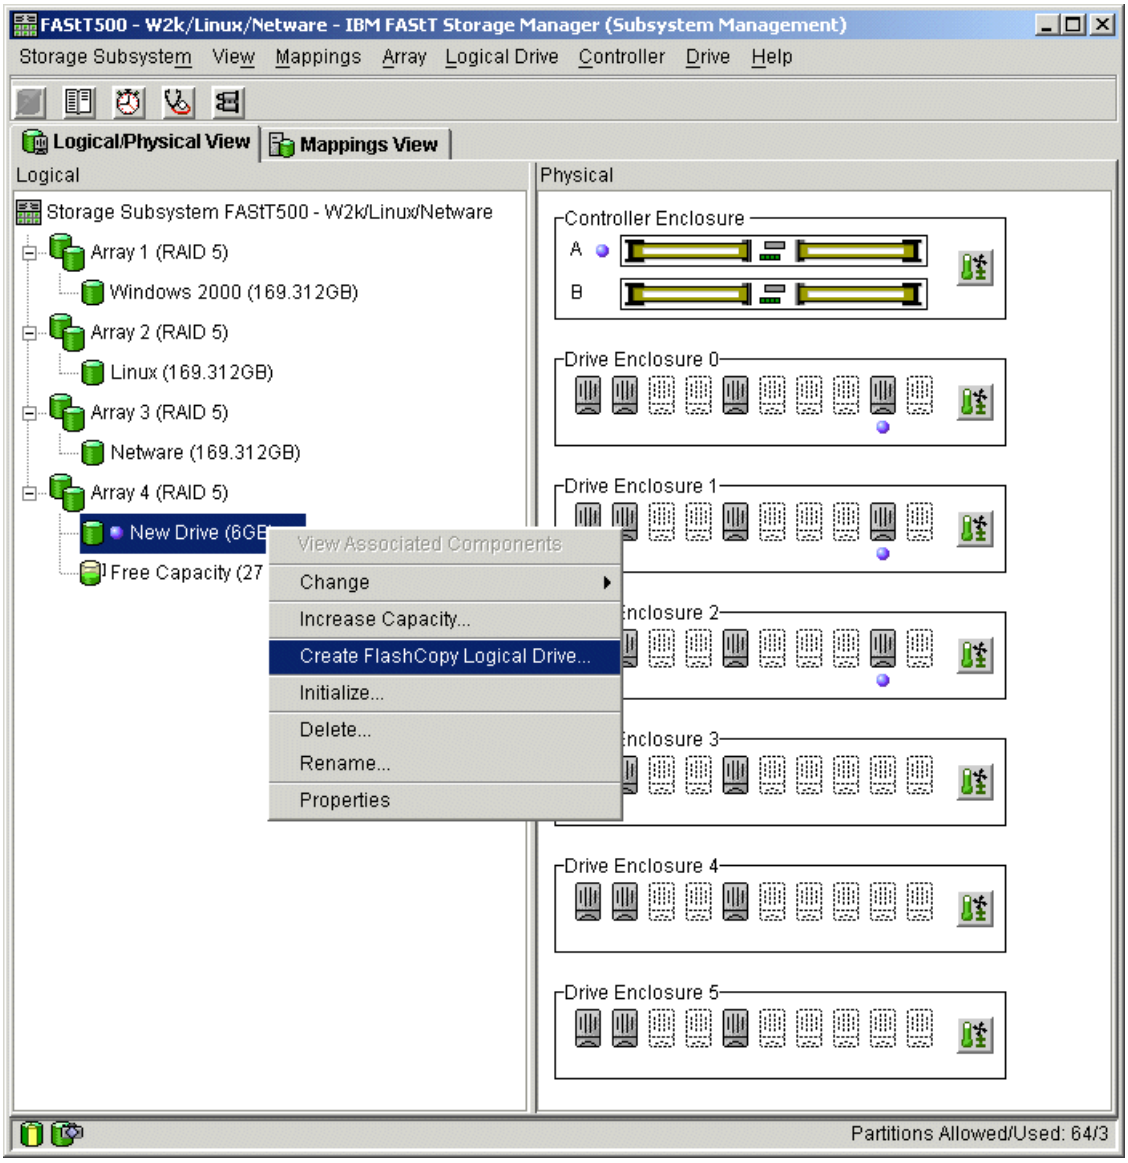

Figure B-6 Storage Manager 8 Subsystems Management window

2. The Create A FlashCopy Logical Drive Wizard - Introduction window is displayed (see Figure B-7):

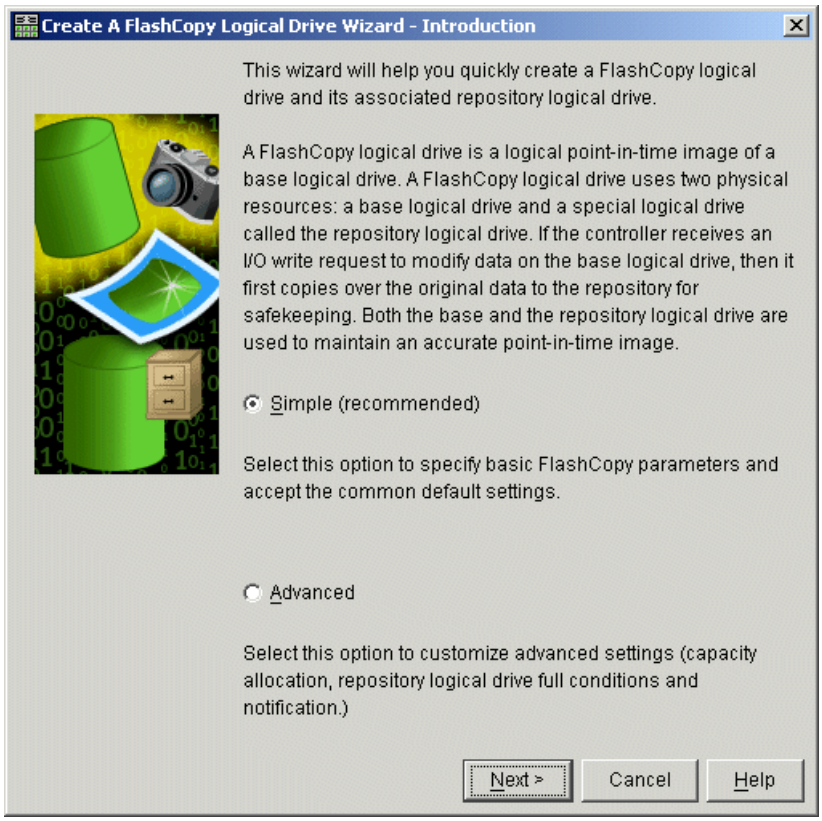

Figure B-7 Create A FlashCopy Logical Drive Wizard - Introduction window

3. For this example we chose the simple method. Select the **Simple** radio button and click **Next**. The Create A FlashCopy Logical Drive Wizard - Specify Names window is displayed (see Figure B-8):

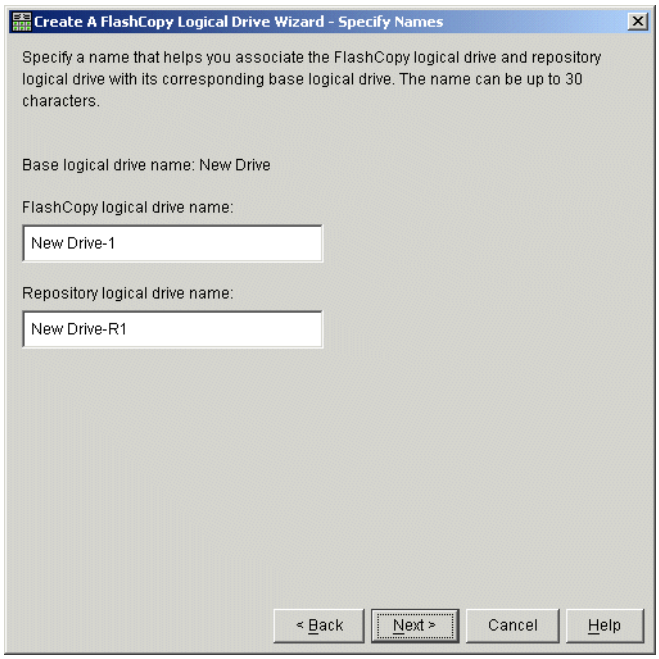

Figure B-8 Create A FlashCopy Logical Drive Wizard - Specify Names window

4. Specify the names of the FlashCopy logical drive and the Repository logical drive, and click **Next** to continue.
5. The Create A FlashCopy Logical Drive Wizard - Specify Repository Capacity window is displayed (see Figure B-9):

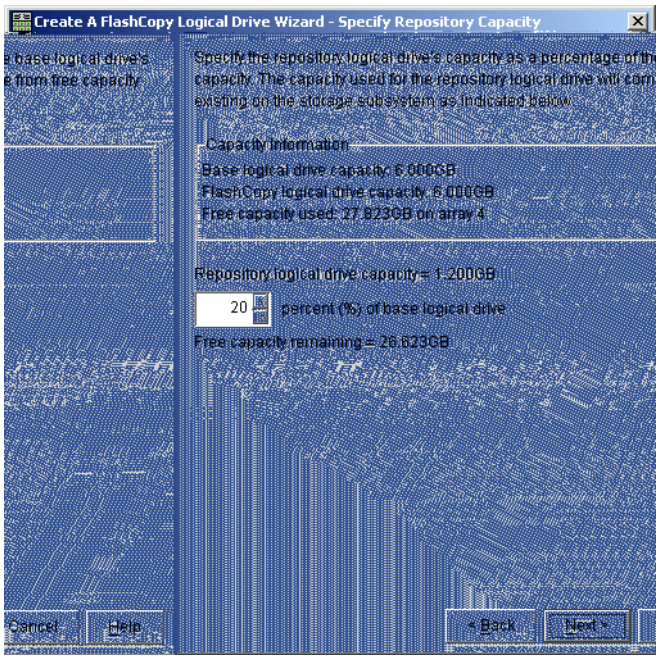

Figure B-9 Create A FlashCopy Logical Drive Wizard - Specify Repository Capacity window

6. Specify how large (in comparison to the original drive) the repository logical drive should be by setting the percentage box to the desired value, then click **Next** to continue.
7. The Create A FlashCopy Logical Drive Wizard - Preview window is displayed (see Figure B-10):

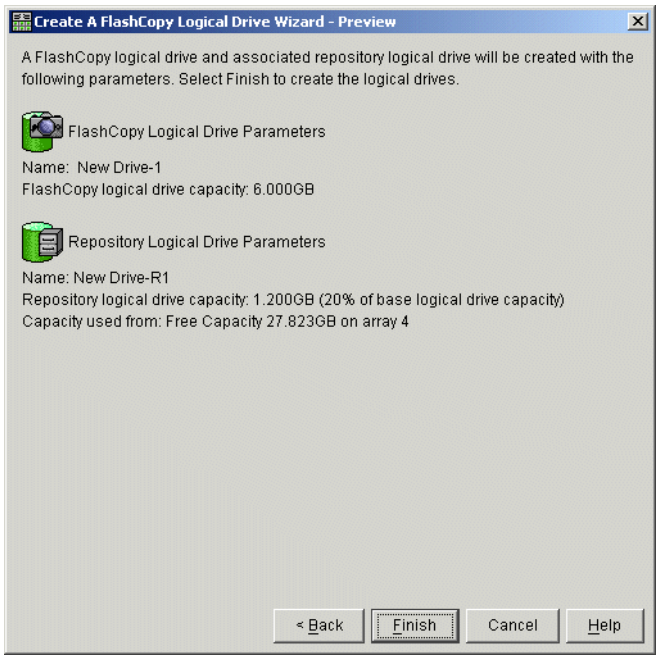

Figure B-10 Create A FlashCopy Logical Drive Wizard - Preview window

8. Check that everything is correct according to your selections and click **Finish** to create the FlashCopy.
9. The Create A FlashCopy Logical Drive Wizard - Completed window will follow, as shown in Figure B-11. Click **OK**.

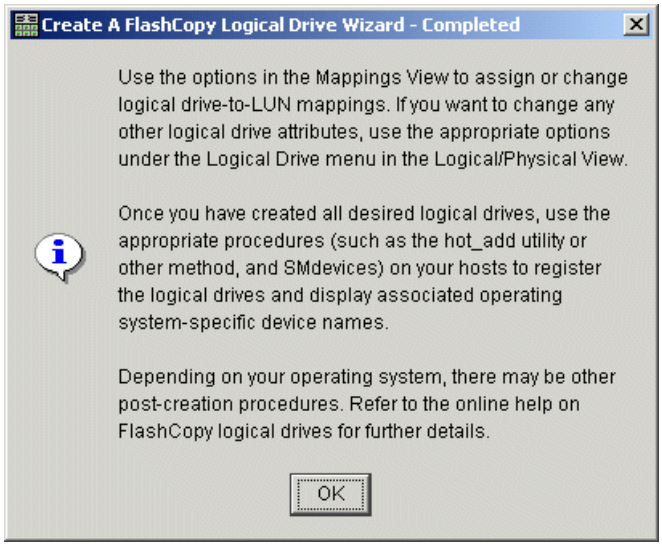

Figure B-11 Create A FlashCopy Logical Drive Wizard - Completed window

The FlashCopy logical drive has now been created.

Deleting a FlashCopy logical drive is as simple as right-clicking the FlashCopy logical drive and selecting **Delete**. For automated management of FlashCopy, it is also possible to create and delete FlashCopy logical drives using the FAStT Storage Manager script editor. You can create scripts to perform these tasks and use the Windows Task Scheduler to schedule them to occur at specific times.

B.2.2 Assigning a FlashCopy logical drive using storage partitioning

Once a FlashCopy logical drive has been created, it is a simple task to assign the disk to a server so that it can be accessed. You can do so using the storage partitioning GUI available in Storage Manager 8. The steps used with the previous version of Storage Manager are also used in the new version. The main improvements are in the way the GUI has been made easier to use, with simpler access to features.

The new GUI is shown in Figure B-12. As you can see, a Mapping View tab has been added to the Subsystem Management window, giving direct access to storage partitioning. The GUI has also been changed to provide a tree view of all host groups and hosts with their assigned logical drives. This is much easier to use then the previous version.

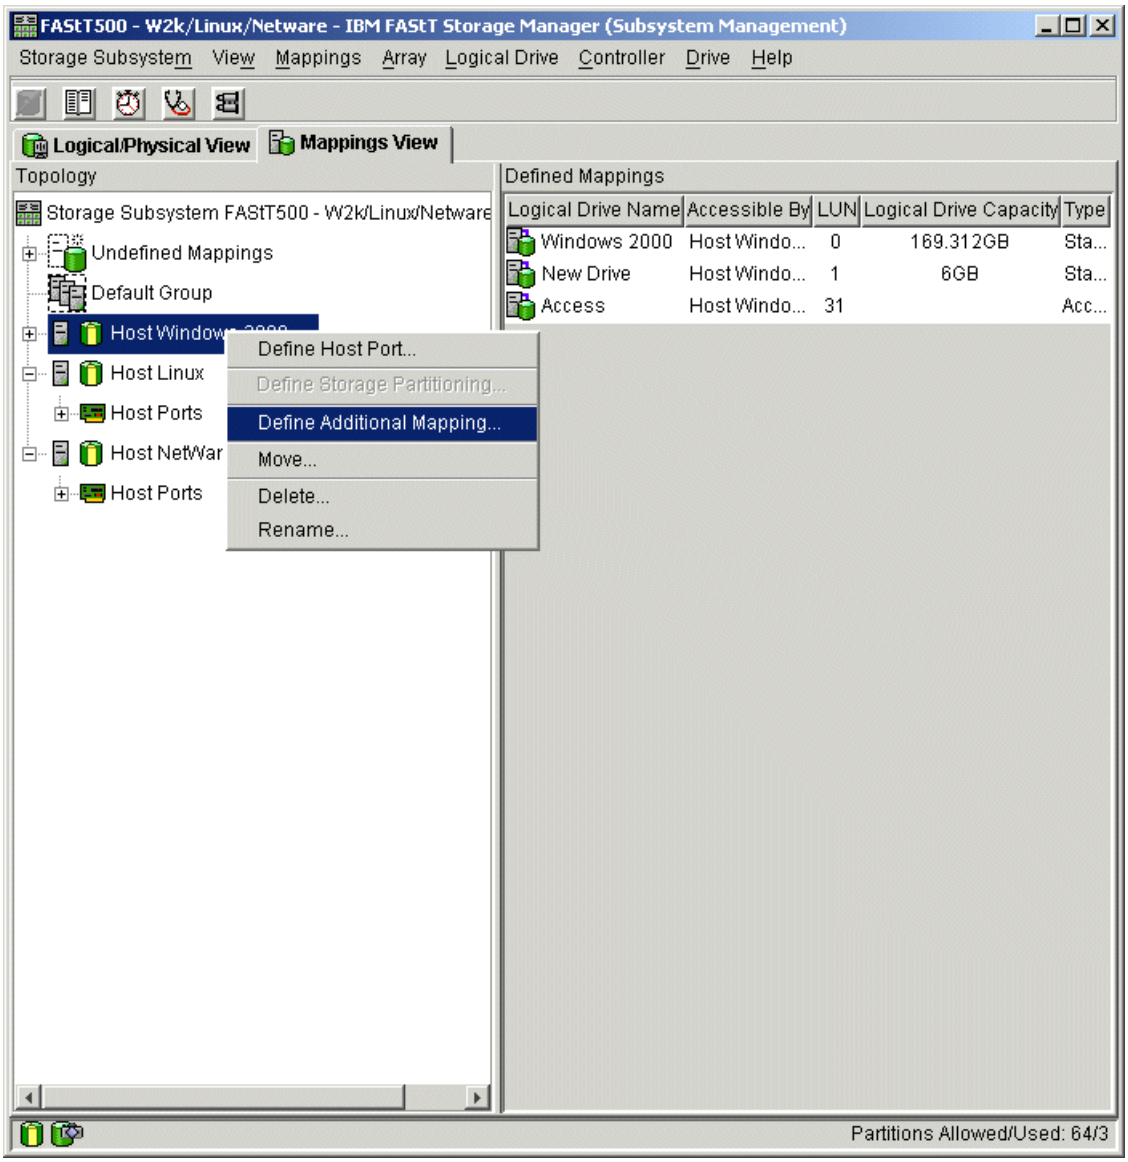

Figure B-12 Storage partitioning in Storage Manager 8

B.2.3 Increasing the capacity of a logical drive

In previous versions of Storage Manager, you had the ability to dynamically increase the size of an array while its logical drives are still being used. You also had the ability to change the RAID level used in the array and could even change the stripe size for a logical drive, again, while the disk was in use.

Now, in the new version of Storage Manager, you also have the ability to increase the size of a logical drive while the subsystem is active. Therefore, using Windows 2000 dynamic disk capability, for example, you can extend a partition without rebooting the server, expanding the usable storage capacity available to your server without downtime.

Note: The ability to increase the size of a logical drive is operating system specific, so please read the online help for each operating system attached to your xSeries SAN.

To perform this feat, you need free space in the array in which the logical drive resides. If your array is full, first increase the size of the array by assigning new disks to it.

Once you have read and understood the specific requirements for the relevant operating system, it is as simple as right-clicking the logical drive, and selecting **Increase capacity**. You then specify the new size you wish to make the logical drive and click **OK**.

The logical drive will automatically start initializing to the new size. You can then go into Windows Disk Management and increase the size of the dynamic partition.

B.3 SANavigator

As storage area networks grow both in size and complexity, it becomes more important to have a centralized management tool to manage your SAN environment in order to keep the cost of ownership to a minimum. In this section we discuss the key functionality of a management tool, called SANavigator.

Whle product plans are not finalized, we believe that a scaled-down version of SANavigator may be bundled with Storage Manager 8. The full version of SANavigator may be purchased from SANavigator.

For more information on SANavigator, visit:

http://www.pc.ibm.com/qtechinfo/MIGR-40262.html

SANavigator consists of two components:

- Server

  This is installed onto a management server or workstation that is permanently switched on. This portion of the software monitors the SAN and performs discoveries for new devices, and so on.

- Client

  This is the GUI for SANavigator, and is installed on administrator workstations or roaming laptop computers to perform remote administration and monitoring.

The features of the scaled-down version of SANavigator include:

- Discovery of up to 32 switch ports

  The ability to manage up to 32 switch ports within SANavigator.

- SAN Monitoring Physical Map

  The ability to monitor the physical connection between SAN devices, and device status (see Figure B-13):

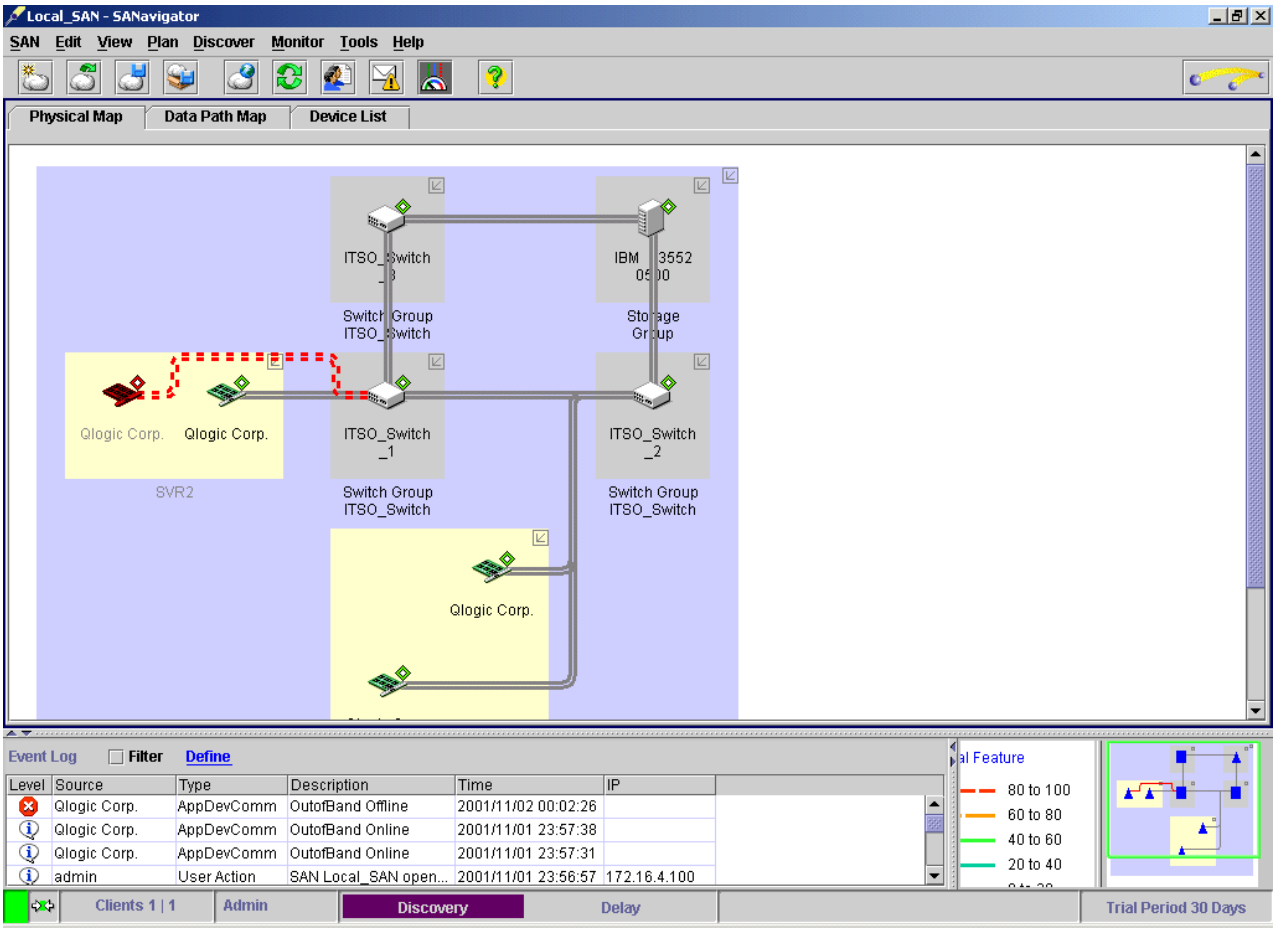

Figure B-13 SANavigator Physical Map with down HBA

- Event Log

  An event log displaying all events that occur during a specified period.

- Event Notification

  The ability to configure alerts to be sent by e-mail or SNMP.

- Device List

  A listing of all devices attached to the SAN (see Figure B-14):

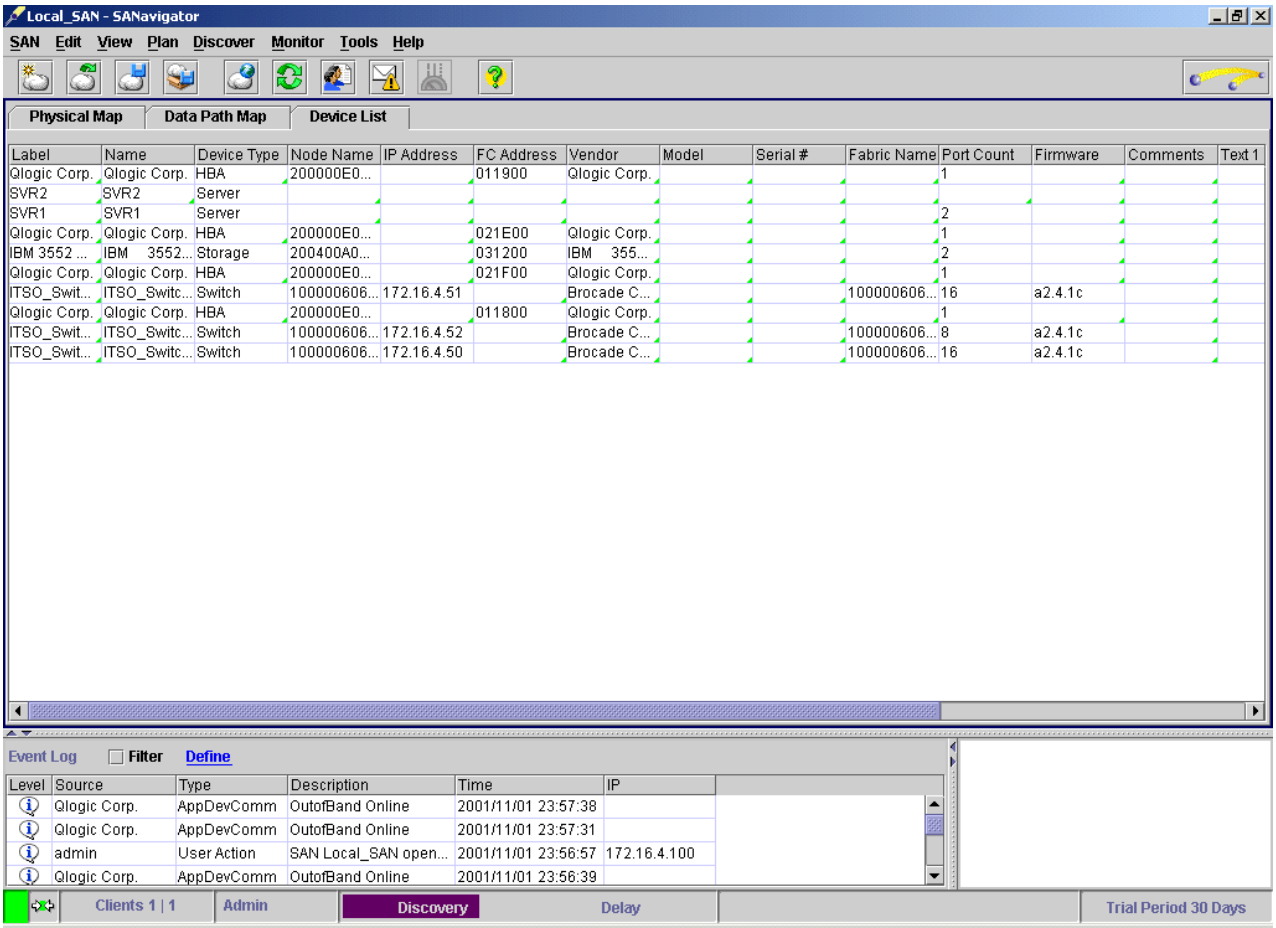

Figure B-14 SANavigator Device List window

The full version of SANavigator adds these features:

- Discovery of more than 32 switch ports

  The ability to manage more than 32 switch ports within SANavigator, licensed per port.

- Performance Monitoring
  - Latency information graphs
  - Switch performance graphs
  - Port performance graphs
  - Performance shown on a physical map
- SAN Planning

  The ability to plan a new SAN or expansion of an existing one.

Special notices

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 other companies:

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

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

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

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

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

UNIX is a registered trademark in the United States and other countries licensed exclusively

© Copyright IBM Corp. 2002

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.